

---

# A Regret-Variance Trade-Off in Online Learning

---

**Dirk van der Hoeven**  
dirk@dirkvanderhoeven.com  
Dept. of Computer Science  
Università degli Studi di Milano, Italy

**Nikita Zhivotovskiy**  
zhivotovskiy@berkeley.edu  
Dept. of Statistics  
University of California, Berkeley

**Nicolò Cesa-Bianchi**  
nicolo.cesa-bianchi@unimi.it  
Dept. of Computer Science  
Università degli Studi di Milano, Italy

## Abstract

We consider prediction with expert advice for strongly convex and bounded losses, and investigate trade-offs between regret and “variance” (i.e., squared difference of learner’s predictions and best expert predictions). With  $K$  experts, the Exponentially Weighted Average (EWA) algorithm is known to achieve  $O(\log K)$  regret. We prove that a variant of EWA either achieves a *negative* regret (i.e., the algorithm outperforms the best expert), or guarantees a  $O(\log K)$  bound on *both* variance and regret. Building on this result, we show several examples of how variance of predictions can be exploited in learning. In the online to batch analysis, we show that a large empirical variance allows to stop the online to batch conversion early and outperform the risk of the best predictor in the class. We also recover the optimal rate of model selection aggregation when we do not consider early stopping. In online prediction with corrupted losses, we show that the effect of corruption on the regret can be compensated by a large variance. In online selective sampling, we design an algorithm that samples less when the variance is large, while guaranteeing the optimal regret bound in expectation. In online learning with abstention, we use a similar term as the variance to derive the first high-probability  $O(\log K)$  regret bound in this setting. Finally, we extend our results to the setting of online linear regression.

## 1 Introduction

In the online learning protocol, the learner interacts with an unknown environment in a sequence of rounds. In each round  $t = 1, 2, \dots, T$  the learner makes a prediction  $\hat{y}_t \in [-\frac{1}{2}M, \frac{1}{2}M]$  and suffers loss  $\ell_t(\hat{y}_t)$ , where  $\ell_1, \ell_2, \dots$  is a sequence of differentiable loss functions unknown to the learner. The goal is to achieve small regret  $\mathcal{R}_T$ , defined by  $\mathcal{R}_T = \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*))$  where  $y_t^*$  are the predictions of some reference forecaster. We consider two special cases of online learning. In prediction with expert advice, the learner receives the predictions  $y_t(1), \dots, y_t(K) \in [-\frac{1}{2}M, \frac{1}{2}M]$  of  $K$  experts at the beginning of each round  $t$ . The learner then predicts with a convex combination  $\hat{y}_t = \sum_{i=1}^K p_t(i)y_t(i)$  of the experts’ predictions, where  $\mathbf{p}_t = (p_t(1), \dots, p_t(K))$  is a probability distribution over experts updated after each round. The goal in this setting is to have small regret with respect to the predictions  $y_t^* = y_t(i^*)$  of any fixed expert  $i^*$ . In online linear regression, the learner has access to a feature vector  $\mathbf{x}_t \in \mathbb{R}^d$  at each round. This is used to compute predictions  $\hat{y}_t = \langle \mathbf{w}_t, \mathbf{x}_t \rangle$ , where  $\mathbf{w}_t \in \mathbb{R}^d$  is a parameter vector to be updated at the end of the round. The goal

is to have small regret with respect to the predictions  $y_t^* = \langle \mathbf{u}, \mathbf{x}_t \rangle$  of any fixed linear forecaster  $\mathbf{u} \in \mathbb{R}^d$  such that  $\|\mathbf{u}\|_2 \leq D$  and  $|y_t^*| \leq \frac{1}{2}M$  for all  $t$ .

In both settings, we assume that the losses  $\ell_t$  are  $\mu$ -strongly convex, which is to say that there exists  $\mu > 0$  such that for all  $y, x \in [-\frac{1}{2}M, \frac{1}{2}M]$ ,

$$\ell_t(y) - \ell_t(x) \leq (y - x)\ell'_t(y) - \frac{\mu}{2}(x - y)^2 \quad t = 1, \dots, T, \quad (1)$$

where we write  $\ell'_t$  to denote the derivative of  $\ell_t$ . Without loss of generality, throughout the paper we assume that  $\mu \leq 2$ .

A typical example of a strongly convex loss is the squared loss  $\ell_t(y) = (y - y_t)^2$  with  $\mu = 2$ . While standard online learning algorithms, such as, for example, Exponentially Weighted Average (EWA) [Vovk, 1990, Littlestone and Warmuth, 1994], are directly applied to the (mixable) losses  $\ell_t$ , we take a different approach. Our algorithms provide bounds on the linearized regret  $\tilde{\mathcal{R}}_T$  of the form

$$\tilde{\mathcal{R}}_T = \sum_{t=1}^T (\hat{y}_t - y_t^*)\ell'_t(\hat{y}_t) \leq \frac{C_T}{\eta} + B_T + \eta \sum_{t=1}^T (\hat{y}_t - y_t^*)^2. \quad (2)$$

for some  $\eta \in (0, H]$  to be chosen freely by the learner and where  $H, C_T, B_T \geq 0$  are problem-specific parameters. In the expert setting, several algorithms provide such a guarantee: Squint [Koolen and Van Erven, 2015], Adapt-ML-Prod [Gaillard et al., 2014], BOA [Wintenberger, 2017], Squint+C and Squint+L [Mhammedi et al., 2019]. In online linear regression, the MetaGrad algorithm [Van Erven et al., 2021] and its variants by Mhammedi et al. [2019], Wang et al. [2020], Chen et al. [2021] satisfy bounds like (2). Although in both settings the aforementioned algorithms achieve the optimal tuning of  $\eta$  in (2) without any preliminary information, our applications do not require this tuning property. As a consequence, the algorithms we derive in Section 3 are simpler and have slightly better guarantees. More importantly, unlike the aforementioned algorithms, we focus on exploiting the curvature of the loss. In particular, by combining equations (1) and (2), we obtain the following lemma, which features the central inequality of our work.

**Lemma 1.** *Consider a sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses. Suppose that predictions  $\hat{y}_t$  guarantee the second-order bound (2). Then*

$$\mathcal{R}_T \leq \frac{C_T}{\eta} + B_T - \left(\frac{\mu}{2} - \eta\right) \sum_{t=1}^T (\hat{y}_t - y_t^*)^2.$$

Next, we show a prototypical example of the bounds we derive in the rest of this work.

**Example 1.** *Consider the expert setting with the squared loss,  $\ell_t(y) = (y - y_t)^2$ , where  $y, y_t \in [-1, 1]$  for all  $t \geq 1$ . Then the predictions  $\hat{y}_t$  of our algorithm satisfy*

$$\mathcal{R}_T \leq 32 \log(K) - \frac{1}{2} \sum_{t=1}^T (\hat{y}_t - y_t^*)^2. \quad (3)$$

Even though (3) seems relatively inconsequential, we exploit the negative variance term in several applications. A straightforward implication is that (3) recovers the usual *constant* regret bound  $\mathcal{R}_T = O(\log(K))$ . More interestingly, however: if we attain the worst-case performance  $O(\log(K))$ , then the variance is small and bounded by  $O(\log(K))$ , if the variance is large enough, then the regret becomes negative. For the role of negative variance terms in the analysis of statistical and online learning, we refer to Section 2.

**Remark.** *Negative quadratic terms similar to the one in Lemma 1 appear in online convex optimization, for example in the analysis of the online gradient descent [Hazan, 2016, Theorem 3.3]. However, as far as we are aware, in online convex optimization the algorithms are not tuned to obtain a negative term in the regret bounds. In this paper we argue that for online regression, and especially for prediction with expert advice, setting  $\eta < \mu/2$  in Lemma 1 and thus obtaining a negative term can be a powerful tool, which plays a central role in all of our applications. In prediction with expert advice the standard approach to exploit the curvature of the loss is through mixability rather than through the negative quadratic terms. In particular, we show that the regret bound in (3) cannot be achieved by EWA with a fixed learning rate, despite the fact that constant regret is achievable by this algorithm via mixability.*

**Proposition 1 (Informal).** *Consider the setup of online prediction with expert advice and bounded strongly convex losses. The bound (3) cannot be achieved by the standard EWA algorithm.*

For the sake of presentation, we defer the formal description of this result to Section 3.1. The proof of Proposition 1 is motivated by a result of Audibert [2007] on the sub-optimality of online to batch converted EWA in deviation.

**Contributions and Outline.** In Section 4, we show our first three applications of Lemma 1 in the framework of statistical learning. In the expert setting, we show that online to batch conversion may be stopped early if the empirical variance of our predictions is sufficiently large. In particular, in high-variance regimes we may stop early because the excess risk bound is negative with high probability, and when the variance is small, we recover the optimal excess risk bound up to a  $\log \log T$  factor. By exploiting the negative variance term again, we show an optimal high-probability excess risk bound for online to batch conversion of algorithms that satisfy (2). The optimal high-probability excess risk bound (called the optimal rate of model selection aggregation) was previously known to be achieved by several estimators appearing in [Audibert, 2007, Lecué and Mendelson, 2009, Lecué and Rigollet, 2014] whose analyses are specific to the statistical learning setup. Our result can also be seen as a simplification and strengthening of a result by Wintenberger [2017]. We also show a high-probability excess risk bound for online to batch conversion of an online regression algorithm in the bounded setup, which is to say that both the feature vectors and derivatives of the losses are bounded. It was previously shown by Mourtada et al. [2021] that online to batch converted versions of the optimal Vovk-Azoury-Warmuth forecaster [Vovk, 2001, Azoury and Warmuth, 2001] have constant excess risk with constant probability.

In Section 5, we use the negative variance term to counteract corrupted feedback in online learning. Our result complements previous results where losses are assumed to be stochastic and corrupted by an adversary—see, e.g., Lykouris et al. [2018], Zimmert and Seldin [2021], Amir et al. [2020], Ito [2021]. In Section 6 we consider the selective sampling setting [Atlas et al., 1989], where the learner’s goal is to control regret while saving on the number of time the current loss is observed. We show that the optimal bound can be recovered while only observing a fraction of all losses if losses are observed proportionally to the cumulative variance. Finally, in Section 7 we discuss how our ideas are not limited to online learning with strongly convex losses, but may also be applied to online learning with abstention and to online multiclass classification. All bounds in the main text with suppressed constants have detailed statements in the appendix. Before discussing our applications, we introduce some notation and discuss the related work. The algorithms that we use to derive most of our results are introduced in Section 3.

**Notation.** We use the standard  $O(\cdot)$  to suppress constants and use  $\tilde{O}(\cdot)$  to suppress constants and  $\log(T)$ -factors. We use  $\log(\cdot)$  to denote the logarithm with base  $e$ . The symbol  $\mathbb{1}[A]$  denotes the indicator of the event  $A$ . For an integer  $K$  we denote  $[K] = \{1, \dots, K\}$ . We use  $p(i) \propto g(i)$  to denote  $p(i) = g(i) / \left( \sum_{i'=1}^K g(i') \right)$  and  $\partial_y \ell(y, Y)$  to denote the derivative of  $\ell$  with respect to  $y$ . The symbol  $I$  denotes the identity matrix whose dimensions are clear from the context. For a set of random variables  $Y, X_1, \dots, X_{t-1}$  we write  $\mathbb{E}_{t-1}[Y] = \mathbb{E}[Y | X_1, \dots, X_{t-1}]$ .

## 2 Related work

**Online learning with losses with curvature.** For a thorough introduction to online learning, we refer the reader to [Cesa-Bianchi and Lugosi, 2006, Hazan, 2016, Orabona, 2019]. Strongly convex losses are a special case of mixable losses. With mixable losses, EWA on a finite set of experts achieves  $O(\log(K))$  regret—see, e.g., [Cesa-Bianchi and Lugosi, 2006, Mhammedi and Williamson, 2018]. Mhammedi and Williamson [2018] observe that—in some cases—regret may be negative for mixable losses, but they do not explore this topic further.

**The role of the negative terms due to curvature in online and statistical learning.** Although bounds of the form (3) are not explicitly present in the literature, the appearance of negative terms is common when proving fast rates in statistical learning with squared loss. In particular, the *empirical star* algorithm of Audibert [2007]—as well as other aggregation algorithms [Lecué and Mendelson, 2009, Lecué and Rigollet, 2014, Wintenberger, 2017]—exploit the curvature of the loss

---

**Algorithm 1:** An algorithm for prediction with expert advice

---

**Input**  $\eta \in (0, \frac{1}{2}]$ ,  $M > 0$  ;  
**Initialize**  $\gamma = \frac{\eta}{M^2}$ ,  $p_1(i) = \frac{1}{K}$  for all  $i$  ;  
**for**  $t = 1, \dots, T$  **do**  
    Receive expert predictions  $y_t(1), \dots, y_t(K)$   
    Predict  $\hat{y}_t = \sum_{i=1}^K p_t(i)y_t(i)$   
    Receive  $g_t$  and  $\kappa_t$   
    Set  $\tilde{\ell}_t(i) = \gamma(y_t(i) - \hat{y}_t)g_t + \kappa_{t-1}(\gamma(y_t(i) - \hat{y}_t)g_t)^2$   
    Set  $p_{t+1}(i) \propto \exp(-\kappa_t \sum_{s=1}^t \tilde{\ell}_s(i))$

---

through the negative term which compensates the variance term. We also refer to some more recent statistical learning results whose analysis is based on the same idea [Mendelson, 2019, Bousquet and Zhivotovskiy, 2021, Puchkin and Zhivotovskiy, 2022, Saad and Blanchard, 2021]. Similarly, in the context of online learning the negative quadratic term appears in Rakhlin and Sridharan [2014], where the so-called *sequential offset Rademacher complexity* is studied. Van Erven et al. [2021] also obtain a bound that is very similar to the bound in Lemma 1 (specifically in the proof of their Theorem 1), but they choose  $\eta$  to match the negative term in their bound. Importantly, in these papers the role of the negative term is only to get the fast rate by compensating the variance term. In contrast, in our case the negative variance terms also appear in the final regret bound and play their role in applications.

**Suboptimality of EWA for prediction with expert advice.** In the setup of prediction with expert advice, the classical Exponentially Weighted Average (EWA) algorithm [Vovk, 1990, Littlestone and Warmuth, 1994] is known to give a constant regret in the case of strongly convex losses. However, despite being optimal in this setting, this algorithm has several known drawbacks. In the case of general losses, EWA does not deliver the second-order bound (2). As a result, unlike more advanced algorithms such as Squint [Koolen and Van Erven, 2015], EWA with a fixed learning rate (and even a decreasing learning rate) cannot adapt to certain benign stochastic environments where, for example, the Bernstein assumption holds [Mourtada and Gaïffas, 2019]. The second source of suboptimality comes from online to batch conversions in the strongly convex case. Although in the statistical setting EWA performs optimally in expectation, it does not do so with high probability [Audibert, 2007]. As a matter of fact, it will be clear from our analysis that this is related to the fact that EWA does not satisfy a bound of the form (3) (see Theorem 1). This problem is one of the motivations behind the work of Wintenberger [2017].

**Exploiting negative regret.** The possibility of getting a negative excess risk has been recently explicitly exploited by Puchkin and Zhivotovskiy [2022, see their Algorithm 3.3 and Theorem 3.4] in the setup of active learning with abstentions. In the context of online to batch conversion of online learning algorithms, a similar idea is exploited in Section 4.1. Moreover, our selective sampling results in Section 6 are of the same flavor.

### 3 Our Algorithms

In the following,  $g_1, g_2, \dots$  are real numbers in a bounded interval. However, in most applications we use  $g_t = \ell'_t(\hat{y}_t)$ . The algorithms in this section are simplified versions of algorithms in the literature. Namely, in Section 3.1 we present a simplified version of Squint [Koolen and Van Erven, 2015] and in Section 3.2 we present a simplified version of MetaGrad [Van Erven et al., 2021]. The simplifications lie in the tuning of the learning rate. Squint and MetaGrad optimize the learning rate online at a small cost, whereas in our applications we only need a fixed learning rate that is known in advance, in which case we do not pay the small cost for optimizing the learning rate. The proofs of the results in this section are postponed to Appendix A.

---

**Algorithm 2:** An algorithm for online linear regression

---

**Input**  $\eta > 0, \sigma > 0, G > 0, Z > 0$  ;  
**Initialize**  $\gamma = \frac{\eta}{G^2}, \mathbf{w}_1 = \mathbf{0}$ , and  $\Sigma_1^{-1} = \frac{1}{\sigma} I$  ;  
**for**  $t = 1, \dots, T$  **do**  
  Receive  $\mathbf{x}_t$   
  Set  $\mathcal{W}_t = \bigcap_{s=1}^t \{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x}_s \rangle| \leq Z\}$   
  Set  $\mathbf{w}_t = \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}_t} (\mathbf{w} - \tilde{\mathbf{w}}_t)^\top \Sigma_t^{-1} (\mathbf{w} - \tilde{\mathbf{w}}_t)$   
  Predict  $\hat{y}_t = \langle \mathbf{w}_t, \mathbf{x}_t \rangle$   
  Receive  $g_t$  and  $\kappa_t$   
  Set  $\mathbf{z}_t = \gamma \mathbf{x}_t g_t$   
  Set  $\Sigma_{t+1}^{-1} = \kappa_t 2 \mathbf{z}_t \mathbf{z}_t^\top + \Sigma_t^{-1}$   
  Set  $\tilde{\mathbf{w}}_{t+1} = \mathbf{w}_t - \mathbf{z}_t \Sigma_{t+1}^{-1}$

---

### 3.1 A Simple Algorithm for Prediction with Expert Advice

Algorithm 1 is a simplified version of Squint [Koolen and Van Erven, 2015]. The  $\kappa_t$  parameter is relevant only to the selective sampling setting, where it is used to control the range of loss estimates. In all other settings we set  $\kappa_t = 1$  for all  $t$ . Note that in round  $t$  both  $\kappa_t$  and  $\kappa_{t-1}$  are used to update the algorithm. This is due to a technicality in the analysis of the algorithm, where a  $\kappa_{t-1} \tilde{\ell}_t(i)$  term appears and we want to use the inequality  $x - x^2 \leq \log(1+x)$  for  $|x| \leq \frac{1}{2}$  (specifically in equation (7)). The regret bound of Algorithm 1 can be found in Lemma 2.

**Lemma 2.** For all  $g_1, \dots, g_T \in [-M, M]$  and  $\kappa_0 = \kappa_1 \geq \kappa_2 \dots \geq \kappa_T$  such that  $\kappa_t \in (0, 1]$ , the predictions  $\hat{y}_t$  of Algorithm 1 run with input  $M$  and  $\eta \in (0, \frac{1}{2}]$  satisfy

$$\sum_{t=1}^T (\hat{y}_t - y_t(i)) g_t \leq \frac{M^2 \log(K)}{\kappa_T \eta} + \eta \sum_{t=1}^T \kappa_{t-1} (\hat{y}_t - y_t^*)^2$$

provided  $\max_i |y_t(i) - y_t^*| \leq M$  for all  $t \geq 1$ .

As an immediate corollary of Lemma 2 and Lemma 1 we have the following regret bound.

**Corollary 1.** Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses such that  $\max_t |\ell'_t| \leq M$ . Provided that  $\max_i |y_t^* - y_t(i)| \leq M$  for all  $t \geq 1$ , the predictions  $\hat{y}_t$  of Algorithm 1 run with inputs  $M$  and  $\eta \in [0, \frac{1}{2})$ ,  $\kappa_t = 1$  for all  $t$ , and feedback  $g_t = \ell'_t(\hat{y}_t)$ , satisfy

$$\mathcal{R}_T \leq \frac{M^2 \log(K)}{\eta} - \left(\frac{\mu}{2} - \eta\right) \sum_{t=1}^T (\hat{y}_t - y_t^*)^2.$$

As an example, let us consider the squared loss, which is 2-strongly convex. In the setup of Example 1, since  $|\ell'_t(\hat{y}_t)| = 2|\hat{y}_t - y_t| \leq 4$ , Algorithm 1 with  $g_t = \ell'_t(\hat{y}_t)$  and  $\eta = \frac{1}{2}$  gives us the regret bound claimed in Example 1, namely  $\mathcal{R}_T \leq 32 \log(K) - \frac{1}{2} \sum_{t=1}^T (\hat{y}_t - y_t^*)^2$ . Our next result is a formal version of Proposition 1 saying that the above regret bound cannot be achieved by the standard EWA algorithm. For standard notation and explicit details on this algorithm we refer to Appendix E.

**Theorem 1.** Consider the squared loss and two experts  $y_t(1) = 0$  and  $y_t(2) = 1$  for all  $t \geq 1$ . Let  $\hat{y}_t^{EWA}$  be the EWA predictions. There is a sequence  $y_1, y_2, \dots$  such that  $y_t \in [0, 1]$ ,  $t \geq 1$  and, for large enough  $T$ , the regret of EWA with  $\eta = \frac{1}{2}$  satisfies  $-12 \log T \leq \mathcal{R}_T \leq 2 \log 2$  and, at the same time,

$$\sum_{t=1}^T (\hat{y}_t^{EWA} - y_t^*)^2 \geq T/2.$$

### 3.2 A Simple Algorithm for Online Regression

In the following we use  $y_t(\mathbf{u}) = \langle \mathbf{u}, \mathbf{x}_t \rangle$ . We prove a regret bound for Algorithm 2, which is a simplified version of MetaGrad [Van Erven et al., 2021]. The role of the parameter  $Z$  in the algorithm is to ensure the predictions  $\hat{y}_t$  are bounded, which will be important in the statistical learning setting in Section 4. Similarly to Algorithm 1, the  $\kappa_t$  parameter is only used in the selective sampling setting.

**Lemma 3.** For all  $g_1, \dots, g_T \in \mathbb{R}$  and  $\kappa_1 \geq \dots \geq \kappa_T \in (0, 1]$ , the predictions  $\hat{y}_t$  of Algorithm 2 run with inputs  $\eta > 0$ ,  $\sigma = D^2$ ,  $G \geq \max_t |g_t|$ , and  $Z > 0$

$$\sum_{t=1}^T (\hat{y}_t - y_t(\mathbf{u}))g_t \leq \frac{dG^2}{2\kappa_T\eta} \log \left( 1 + D^2\eta^2 \left( \max_{t=1, \dots, T} \|\mathbf{x}_t\|_2^2 \right) \frac{T}{d} \right) + \frac{G^2}{2\eta} + \eta \sum_{t=1}^T \kappa_t (\hat{y}_t - y_t(\mathbf{u}))^2,$$

for any  $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{R}^d$ , and for any  $\mathbf{u} \in \mathcal{W}_T \equiv \bigcap_{t=1}^T \{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x}_t \rangle| \leq Z\}$  such that  $\|\mathbf{u}\|_2 \leq D$ .

**Example 2.** Consider the setup of Example 1 and suppose that  $\max_t \|\mathbf{x}_t\|_2, \|\mathbf{u}\|_2 \leq 1$  and  $M = 1$ . We have that  $|\ell'_t(\hat{y}_t)| = |2(\hat{y}_t - y_t)| \leq 4$  and thus, by Lemma 3, an appropriately tuned Algorithm 2 with  $g_t = \ell'_t(\hat{y}_t)$  satisfies (2) with  $C_T = 8 + 8d \log(1 + \eta^2 \frac{T}{d})$  and  $B_T = 0$ . Thus, by Lemma 1, setting  $\eta = \frac{1}{2}$  gives us

$$R_T \leq 16 + 16d \log\left(1 + \frac{T}{4d}\right) - \frac{1}{2} \sum_{t=1}^T (\hat{y}_t - y_t(\mathbf{u}))^2.$$

## 4 Statistical Learning

We discuss an application of our general results in the context of statistical learning where we are interested in the generalization of estimators to unseen samples. A tool often used in converting online learning algorithms to the statistical learning setting is online to batch conversion [Cesa-Bianchi et al., 2004]. Let us recall the setup.

Assume that we are given a family  $\mathcal{F}$  of real-valued functions defined on the instance space  $\mathcal{X}$ . We observe  $T$  i.i.d. observations  $(X_t, Y_t)_{t=1}^T$  distributed according to some unknown distribution  $\mathbb{P}$  on  $\mathcal{X} \times \mathbb{R}$ . Given the loss function  $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}$ , define the *risk*  $R(f)$  of  $f : \mathcal{X} \rightarrow \mathbb{R}$  as  $R(f) = \mathbb{E} \ell(f(X), Y)$ , where the expectation is taken with respect to the joint distribution of  $X$  and  $Y$ . We are interested in bounding the *excess risk*  $R(\hat{f}) - \inf_{f \in \mathcal{F}} R(f)$ , where  $\hat{f}$  is constructed

based on the sample  $(X_t, Y_t)_{t=1}^T$ . Assume that there is a sequence of predictors  $\hat{f}_1, \dots, \hat{f}_T$  trained in an online manner using  $(X_t, Y_t)_{t=1}^T$  (that is,  $\hat{f}_k$  depends on  $(X_t, Y_t)_{t=1}^{k-1}$ ) such that almost surely  $\sum_{t=1}^T \left( \ell(\hat{f}_t(X_t), Y_t) - \ell(f^*(X_t), Y_t) \right) \leq R_T$ , where  $R_T$  is non-random. In this case, a standard

online to batch conversion approach gives an in-expectation excess risk bound  $\mathbb{E} R\left(\frac{1}{T} \sum_{t=1}^T \hat{f}_t\right) - \inf_{f \in \mathcal{F}} R(f) \leq \frac{R_T}{T}$ , for any loss convex in its first argument and where the expectation is taken with

respect to the learning sample  $(X_t, Y_t)_{t=1}^T$ . However, getting a high-probability version of this result is a known challenge if one wants to get the fast rate  $O\left(\frac{1}{T}\right)$ . A standard way of proving a high-probability result is to apply Freedman's inequality for martingales [Kakade and Tewari, 2008] that leads to a variance term scaling as  $O\left(\frac{1}{\sqrt{T}}\right)$ . For example, Audibert [2007] showed that this is the case if one wants to prove a high-probability excess risk bounds based on EWA. A way to handle the variance term in Freedman's inequality is by exploiting the Bernstein assumption as in Kakade and Tewari [2008]. Unfortunately, this assumption is not necessarily satisfied by the stochastic environments we are considering. The main idea in this section is to use the negative term from Lemma 1 to cancel out this variance term appearing due to Freedman's inequality<sup>1</sup>. We use the following notation when applying the online algorithms in the statistical setting:  $\ell_t(\cdot) = \ell(\cdot, Y_t)$  and  $y_t(f) = f(X_t)$ .

### 4.1 Statistical Learning: Model Selection Aggregation

In this section, we discuss the application of our results to the model selection (MS) aggregation. This setup was introduced by Nemirovski [2000] and further studied by Tsybakov [2003] and by Audibert [2007], Lecué and Mendelson [2009], Lecué and Rigollet [2014], Wintenberger [2017], Mourtada et al. [2021] among other works. We introduce some notation. Assume that we are given a finite dictionary  $\mathcal{F} = \{f_1, \dots, f_K\}$  of real-valued functions defined on the instance space  $\mathcal{X}$ . We observe  $T$  i.i.d. observations  $(X_t, Y_t)_{t=1}^T$  distributed according to some unknown distribution  $\mathbb{P}$  on  $\mathcal{X} \times \mathbb{R}$ .

<sup>1</sup>We remark that Wintenberger [2017] uses a similar but technically more involved idea to compensate the variance of predictions using the term appearing because of the curvature of the loss.

---

**Algorithm 3:** Early Stopping online to batch for Model Selection Aggregation

---

**Input**  $T, M, \eta$ , stopping threshold  $S$  ;

**Initialize**  $S = 0$ , provide  $\eta$ , and  $M$  as input for Algorithm 1 ;

**while**  $S < T$  and  $S > \frac{\mu}{8} \min_{f \in \mathcal{F}} \sum_{t=1}^S (\hat{y}_t - f(X_t))^2$  **do**  
  Receive  $X_t$  and send  $f_1(X_t), \dots, f_K(X_t)$  as expert predictions to Algorithm 1  
  Receive  $\mathbf{p}_t$  and  $\hat{y}_t = \sum_{i=1}^K p_t(i) f_i(X_t)$  from Algorithm 1  
  Predict  $\hat{y}_t$  and receive  $\ell_t$   
  Send  $g_t = \ell'_t(\hat{y}_t)$  and  $\kappa_t = 1$  to Algorithm 1  
  Set  $S = S + 1$   
**Output**  $\hat{f} = \frac{1}{S} \sum_{t=1}^S \sum_{i=1}^K p_t(i) f_i$  ;

---

Given the loss function  $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}$ , define the *risk*  $R(f)$  of  $f : \mathcal{X} \rightarrow \mathbb{R}$  as  $R(f) = \mathbb{E} \ell(f(X), Y)$ , where the expectation is taken with respect to the joint distribution of  $X$  and  $Y$ . In the model selection aggregation, one is interested in constructing an estimator  $\hat{f}$  based on the random sample  $(X_t, Y_t)_{t=1}^T$  such that, with probability at least  $1 - \delta$ ,

$$R(\hat{f}) - \min_{f \in \mathcal{F}} R(f) = O\left(\frac{\log(K) + \log(1/\delta)}{T}\right), \quad (4)$$

under appropriate boundedness and curvature assumptions on the loss function  $\ell$ . Analogously to Tsybakov [2003], the bound of the form (4) will be called the *optimal rate of aggregation*. We make use of a variant of online to batch conversion [Cesa-Bianchi et al., 2004] where we stop the procedure early if the empirical variance of predictions is sufficiently large. We sketch the idea. Let  $S$  be the number of samples we have used before we terminated the procedure. We use Algorithm 1 as our aggregation procedure and use  $\hat{f} = \frac{1}{S} \sum_{t=1}^S \sum_{i=1}^K p_t(i) f_i$ . By Jensen's inequality we have  $R(\hat{f}) \leq \frac{1}{S} \sum_{t=1}^S \mathbb{E}_{t-1} \left[ \ell_t \left( \sum_{i=1}^K p_t(i) f_i(X_t) \right) \right]$ . To motivate stopping early, observe that if the empirical variance in Lemma 1 is sufficiently large, we may conclude that the excess risk is negative and we have outperformed the best  $f \in \mathcal{F}$ . The result can be found in Theorem 2 below, whose proof is implied by Theorem 8 in Appendix B.

**Theorem 2.** *Suppose that for all  $f \in \mathcal{F}$   $|f(X)| \leq \frac{1}{2}$  almost surely, that  $|\partial_y \ell(y, Y)| \leq 1$  almost surely for all  $y$  such that  $|y| \leq \frac{1}{2}$ , and that  $\ell$  is  $\mu$ -strongly convex in its first argument. Then, with probability at least  $1 - \delta$ , Algorithm 3 with input parameters  $T, S = O\left(\frac{\log(K) + \log(\log(T)/\delta)}{\mu}\right)$ ,  $\eta = \frac{\mu}{8}$ , and  $M = 1$  satisfies*

$$R(\hat{f}) \leq \begin{cases} \min_{f \in \mathcal{F}} R(f) & \text{if } S < T \\ \min_{f \in \mathcal{F}} R(f) + O\left(\frac{\log(K) + \log(\log(T)/\delta)}{\mu T}\right) & \text{if } S = T, \end{cases}$$

where  $S$  is the number of steps of Algorithm 3.

When Algorithm 3 terminates at step  $S = T$ , we recover the optimal high probability bound for model selection aggregation (4) up to an additive  $\log \log T$  term. However, when  $S = T$  our bound tells us slightly more, because we know that for all  $t' < T$ ,  $\min_{f \in \mathcal{F}} \sum_{t=1}^{t'} (\hat{y}_t - f(X_t))^2 = O(\log K + \log \log T)$ , which means that on the sequence  $(X_t, Y_t)_{t=1}^T$  our predictions  $\hat{y}_t$  are essentially following the prediction of the currently best expert at each round.

In the special case where we are solely interested in the best possible performance of the online to batch conversion of Algorithm 1, we can remove the  $\log \log T$  term appearing in the previous bound. We remark that, apart from the work of Wintenberger [2017], no known analysis based on the online to batch conversion achieved the optimal rate of aggregation (4). We also believe that our analysis is simpler than for previously known algorithms. The result can be found in Theorem 3 below, whose result is implied by Theorem 9 in Appendix B.

**Theorem 3.** *Suppose that for all  $f \in \mathcal{F}$   $|f(X)| \leq \frac{1}{2}$  almost surely, that  $|\partial_y \ell(y, Y)| \leq 1$  almost surely for all  $y$  such that  $|y| \leq \frac{1}{2}$ , and that  $\ell$  is  $\mu$ -strongly convex in its first argument. Then, with probability at least  $1 - \delta$ , Algorithm 3 with input parameters  $T, \eta = \frac{\mu}{4}, S = \infty$ , and  $M = 1$ ,*

guarantees

$$R(\hat{f}) - \min_{f \in \mathcal{F}} R(f) = O\left(\frac{\log(K) + \log(1/\delta)}{\mu T}\right).$$

## 4.2 Statistical Learning: Regression

We consider the statistical learning setting where one has access to  $T$  i.i.d. samples of pairs  $(X_t, Y_t) \in \mathbb{R}^d \times \mathbb{R}$ . In this case, we consider  $\mathcal{F} \subseteq \{\mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle : \mathbf{w} \in \mathbb{R}^d\}$ . For  $\mathbf{w} \in \mathbb{R}^d$  we define the risk as  $R(\mathbf{w}) = \mathbb{E}[\ell(\langle \mathbf{w}, X \rangle, Y)]$ . We assume that  $\ell$  is  $\mu$ -strongly convex in its first argument.

To the best of our knowledge there are no known high probability excess risk bounds in regression based on online to batch conversions with convergence rate  $O\left(\frac{d \log(T)}{T}\right)$ . We provide such a result in the bounded setup where the feature vectors, derivatives of the losses, and the norm of the reference vector are bounded. Similarly to before, for a result that holds with high probability, one needs to control the cumulative variance of our prediction. For standard online learning algorithms the control of the variance may prove troublesome. For example, Mourtada et al. [2021] showed that a version of Vovk-Azoury-Warmuth forecaster [Vovk, 2001, Azoury and Warmuth, 2001] may have a  $O(1)$  excess risk bound with constant probability, whereas in expectation the Vovk-Azoury-Warmuth forecaster guarantees a  $O\left(\frac{d \log(T)}{T}\right)$  excess risk bound. Instead, we leverage the negative empirical variance of Lemma 1 to control the variance of the online to batch conversion, leading to the following excess risk bound, whose result is implied by Theorem 10 in Appendix B.

**Theorem 4.** *Suppose that  $\|X\|_2 \leq 1$  and  $\sup_{y \in [-1, 1]} |\partial_y \ell(y, Y)| \leq 1$  almost surely,  $\|\mathbf{w}\|_2 \leq 1$ , and that  $\ell$  is  $\mu$ -strongly convex in its first argument. Then, with probability at least  $1 - \delta$ ,*

$$R\left(\frac{1}{T} \sum_{t=1}^T \mathbf{w}_t\right) - R(\mathbf{w}) = O\left(\frac{d \log(T) + \log(1/\delta)}{\mu T}\right),$$

where  $\mathbf{w}_t$  are given by Algorithm 2 with  $\eta = \frac{\mu}{4}$ ,  $\sigma = 1$ ,  $Z = 1$ ,  $G = 1$ ,  $\kappa_t = 1$ , and feedback  $g_t = \ell'_t(\langle \mathbf{w}_t, X_t \rangle)$  for  $t = 1, \dots, T$ .

## 5 Corrupted Feedback

In this section, we study a setting where the loss derivatives  $\ell'_t(\hat{y}_t)$  observed by the learner at each round  $t$  may be adversarially corrupted by unknown additive constants  $c_t$ , and we are interested in the best possible dependence on  $c_1, \dots, c_T$  in the regret bound. To better explain our setting, we start with the following example.

**Example 3.** *In the online regression setting, suppose that  $\ell_t(\hat{y}_t) = (\hat{y}_t - y_t)^2$  for all  $t$ , but the learner observes corrupted outcomes  $y_t - c_t/2$ . Hence, the squared loss derivative computed by the learner is  $2(\hat{y}_t - y_t + c_t/2) = \ell'_t(\hat{y}_t) + c_t$ , which can be handled by the algorithms developed in this section.*

Several variants of this setting have been studied in prior work, see for example [Lykouris et al., 2018, Amir et al., 2020, Zimmert and Seldin, 2021, Ito, 2021] and the references therein. The main difference between our setting and these previous settings is that we assume our losses to be strongly convex and the environment is not necessarily stochastic. Although the results in this section are rather straightforward corollaries of our bounds, we believe that it is instructive to provide some explicit results. All proofs of the results in this section are postponed to Appendix C.

Our first result shows the performance of Algorithm 1 in the setup with corrupted gradients. The proof follows from observing that  $\ell'_t(\hat{y}_t)(\hat{y}_t - y_t^*) = (\ell'_t(\hat{y}_t) + c_t)(\hat{y}_t - y_t^*) - c_t(\hat{y}_t - y_t^*)$  and that for any  $\lambda > 0$ , the inequality  $|c_t(\hat{y}_t - y_t^*)| \leq \frac{c_t^2}{\lambda} + \frac{\lambda}{4}(\hat{y}_t - y_t^*)^2$  holds. The  $\frac{\lambda}{4}(\hat{y}_t - y_t^*)^2$  term can be compensated for by the negative  $\frac{\mu}{2}(\hat{y}_t - y_t^*)^2$  appearing in Lemma 1, leading to a  $\sum_{t=1}^T c_t^2$  additive term in the regret bound. In particular, our result implies that as long as  $\sum_{t=1}^T c_t^2$  is of order  $O(\log K)$ , the same regret bound as if the losses were not corrupted. The formal statement can be found in Theorem 5 below.

**Theorem 5.** *Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses and corruptions  $c_1, \dots, c_T \in \mathbb{R}$ . Then the predictions  $\hat{y}_t$  of Algorithm 1 run with inputs  $M \geq$*



$\max_t |\ell'_t(\hat{y}_t) + c_t|$ ,  $\eta = \frac{\mu}{4}$ , feedback  $g_t = \ell'_t(\hat{y}_t) + c_t$ , and  $\kappa_t = 1$  satisfy

$$\mathcal{R}_T \leq \frac{8M^2 \log(K)}{\mu} + \sum_{t=1}^T \frac{c_t^2}{\mu} - \frac{\mu}{8} \sum_{t=1}^T (\hat{y}_t - y_t(i^*))^2$$

provided  $\max_i \max_t |\hat{y}_t - y_t(i)| \leq M$ .

Next we prove an analog of Theorem 5 in the online regression setup.

**Theorem 6.** Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses and corruptions  $c_1, \dots, c_T \in \mathbb{R}$ . Then the predictions  $\hat{y}_t$  of Algorithm 2 run with inputs  $\eta = \frac{\mu}{8}$ ,  $\sigma = D^2$ ,  $G \geq \max_t |\ell'_t(\hat{y}_t) + c_t|$ ,  $Z > 0$ , feedback  $g_t = \ell'_t(\hat{y}_t) + c_t$ , and  $\kappa_t = 1$ , satisfy

$$\mathcal{R}_T \leq \frac{4dG^2}{\mu} \log \left( 1 + \frac{TD^2\mu^2 \max_t \|\mathbf{x}_t\|_2^2}{2d} \right) + \frac{4G^2}{\mu} + \sum_{t=1}^T \frac{c_t^2}{\mu} - \frac{\mu}{8} \sum_{t=1}^T (\hat{y}_t - y_t^*)^2,$$

for any  $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{R}^d$ , and for any  $\mathbf{u} \in \mathcal{W}_T \equiv \bigcap_{t=1}^T \{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x}_t \rangle| \leq Z\}$  such that  $\|\mathbf{u}\|_2 \leq D$  and  $y_t^* = \langle \mathbf{u}, \mathbf{x}_t \rangle$  for all  $t \geq 1$ .

## 6 Selective Sampling

We consider a variant of the selective sampling setting—see, e.g., [Atlas et al., 1989, Freund et al., 1997, Cesa-Bianchi et al., 2003, 2006, Orabona and Cesa-Bianchi, 2011]—where the learner has access to the expert predictions (or, equivalently, to feature vectors), but can observe its own loss only upon request. The goal is to trade off the number of loss requests with regret guarantees. We show that if the variance is high, with only a fraction of all losses requested we obtain the same guarantee (in expectation and up to constants) as when all losses are requested.

Let  $o_t = 1$  with probability  $q_t$  and  $o_t = 0$  with probability  $1 - q_t$ . In each round, if  $o_t = 1$  the loss  $\ell_t$  at round  $t$  is requested, and we use the loss estimator  $\frac{o_t}{q_{t-1}} \ell_t$  to update. Note that this is not the importance weighted estimator, as  $\mathbb{E}_{t-1}[\frac{o_t}{q_{t-1}}] = \frac{q_t}{q_{t-1}}$ . The reason for choosing this particular loss estimator is that we have better control of the range of the loss which allows us to tune  $\kappa_t$  in Algorithm 1 accordingly. The probability of requesting a loss is

$$q_t = \min \left\{ 1, \beta \sqrt{\min_i \sum_{s=1}^t (\hat{y}_s - y_s(i))^2} \right\}, \quad (5)$$

where  $\beta > 0$  is chosen by the learner. Our result for the selective sampling setting can be found in Theorem 7 below, whose statement is implied by Theorem 11 in Appendix D. Theorem 7 implies the following: if  $\beta = O(\mu^{-3/2} \log(K))$ , then with only an expected number  $\sum_{t=1}^T q_t$  of loss requests, we obtain (up to constants) the same regret guarantee as we would have obtained if we had requested all losses. With this particular choice of  $\beta$ ,  $q_t < 1$  as soon the bound in Lemma 1 becomes negative. In other words, when the variance is high we only need a fraction of the losses to recover the worst-case optimal regret bound (in expectation). A similar result can be obtained in the regression setting, see Appendix D.1.

**Theorem 7.** Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses. Provided  $\max_i \max_t |\hat{y}_t - y_t(i)| \leq 1$ , the predictions  $\hat{y}_t$  of Algorithm 1 run with inputs  $M = 1 \geq \max_t |\ell'_t(\hat{y}_t)|$ ,  $\eta = \frac{\mu}{4}$ , feedback  $g_t = \frac{o_t}{q_{t-1}} \ell'_t(\hat{y}_t)$ , and  $\kappa_t = q_t$  satisfy

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] = O \left( \frac{\log(K)}{\mu} + \frac{\log(K)^2}{\mu^3 \beta^2} \right).$$

## 7 Further Extensions

In Appendix G we present another application of Lemma 1. Namely, we show that we may restart Algorithm 1 for free whenever the regret becomes negative. This gives us a regret bound where we compete with a new expert after each restart, which is a stronger notion of regret than when we compete with a fixed expert in all rounds.

The ideas of Lemma 1 also naturally extend beyond online learning with strongly convex losses. We discuss one such extension to the online prediction with expert advice setting. The online prediction with abstention setting was introduced by Neu and Zhivotovskiy [2020] and proceeds as follows. In each round  $t = 1, \dots, T$  the learner receives expert predictions  $y_t(i) \in \{-1, 1\}$ ,  $i = 1, \dots, K$  and the learner can then either predict  $\tilde{y}_t \in \{-1, 1\}$  or abstain from prediction. If the learner predicts with  $\tilde{y}_t$ , the learner suffers the binary loss  $\ell_t(\tilde{y}_t) = \mathbb{1}[\tilde{y}_t \neq y_t]$ , where  $y_t \in \{-1, 1\}$ . If the learner abstains from prediction, the learner suffers abstention cost  $\rho \in [0, \frac{1}{2})$ . Let  $a_t = 1$  for prediction and  $a_t = 0$  for abstention. The total loss of the learner is therefore equal to

$$\sum_{t=1}^T (a_t \mathbb{1}[\tilde{y}_t \neq y_t] + (1 - a_t)\rho).$$

Assume that the prediction strategy is random in a sense that  $a_t$  are Bernoulli random variables whose means might depend on previous observations. The work of Neu and Zhivotovskiy [2020] shows that there is a randomized prediction strategy such that for any data generating mechanism it holds that

$$\mathbb{E} \left[ \sum_{t=1}^T (a_t \mathbb{1}[\tilde{y}_t \neq y_t] + (1 - a_t)\rho) \right] - \sum_{t=1}^T \mathbb{1}[y_t^* \neq y_t] = O \left( \frac{\log K}{1 - 2\rho} \right), \quad (6)$$

independently of  $T$ , where the expectation is taken with respect to the randomness of  $a_t$ ; here  $y_t^* = y_t(i^*)$  and  $i^* = \operatorname{argmin}_i \sum_{t=1}^T \mathbb{1}[y_t(i) \neq y_t]$ . Although it was shown that the randomization is necessary to achieve the regret bound (6), it is unclear if the same regret bound can be achieved with high probability with respect to the randomization of the algorithm. In Appendix F, we answer this question using the techniques we developed and provide a randomized algorithm such that, with probability at least  $1 - \delta$ ,

$$\sum_{t=1}^T (a_t \mathbb{1}[\tilde{y}_t \neq y_t] + (1 - a_t)\rho) - \sum_{t=1}^T \mathbb{1}[y_t^* \neq y_t] = O \left( \frac{\log K + \log(1/\delta)}{1 - 2\rho} \right).$$

In Appendix F we prove Lemma 7, which is the analog of Lemma 1 for the abstention setting. The equivalent of the negative term in Lemma 1 is used to compensate for the variance of the high-probability statement, which allows us to recover the above bound. This also implies that our other applications of Lemma 1 can be exported to the online learning with abstention setting.

## Acknowledgments and Disclosure of Funding

Dirk van der Hoeven and Nicolò Cesa-Bianchi gratefully acknowledge partial support from the MIUR PRIN grant Algorithms, Games, and Digital Markets (ALGADIMAR), the EU Horizon 2020 ICT-48 research and innovation action under grant agreement 951847, project ELISE (European Learning and Intelligent Systems Excellence), and the project ‘‘One Health Action Hub: University Task Force for the resilience of territorial ecosystems’’ funded by Università degli Studi di Milano. A significant part of this work was done while Nikita Zhivotovskiy was at ETH, Zürich. During this period, Nikita Zhivotovskiy was funded in part by ETH Foundations of Data Science (ETH-FDS).

## 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.
- Les Atlas, David Cohn, and Richard Ladner. Training connectionist networks with queries and selective sampling. *Advances in Neural Information Processing Systems*, 2, 1989.
- Jean-Yves Audibert. Progressive mixture rules are deviation suboptimal. *Advances in Neural Information Processing Systems*, 2007.
- Katy S Azoury and Manfred K Warmuth. Relative loss bounds for on-line density estimation with the exponential family of distributions. *Machine Learning*, 43(3):211–246, 2001.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *International Conference on Artificial Intelligence and Statistics*, pages 19–26, 2011.

- Olivier Bousquet and Nikita Zhivotovskiy. Fast classification rates without standard margin assumptions. *Information and Inference: A Journal of the IMA*, 10(4):1389–1421, 2021.
- Sébastien Bubeck. Introduction to online optimization. *Lecture Notes*, 2011.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Nicolò Cesa-Bianchi, Alex Conconi, and Claudio Gentile. Learning probabilistic linear-threshold classifiers via selective sampling. In *Learning Theory and Kernel Machines*, pages 373–387. Springer, 2003.
- Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- Nicolo Cesa-Bianchi, Claudio Gentile, Luca Zaniboni, and Manfred Warmuth. Worst-case analysis of selective sampling for linear classification. *Journal of Machine Learning Research*, 7(7), 2006.
- 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.
- 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.
- Tim van Erven, Wouter M. Koolen, and Dirk van der Hoeven. Metagrad: Adaptation using multiple learning rates in online learning. *Journal of Machine Learning Research*, 22(161):1–61, 2021.
- Yoav Freund, H Sebastian Seung, Eli Shamir, and Naftali Tishby. Selective sampling using the query by committee algorithm. *Machine Learning*, 28(2):133–168, 1997.
- Pierre Gaillard, Gilles Stoltz, and Tim Van Erven. A second-order bound with excess losses. In *Conference on Learning Theory*, pages 176–196, 2014.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- Dirk van der Hoeven. Exploiting the surrogate gap in online multiclass classification. *Advances in Neural Information Processing Systems*, 33:9562–9572, 2020.
- Dirk van der Hoeven, Tim van Erven, and Wojciech Kotłowski. The many faces of exponential weights in online learning. In *Conference On Learning Theory*, pages 2067–2092, 2018.
- Shinji Ito. On optimal robustness to adversarial corruption in online decision problems. *Advances in Neural Information Processing Systems*, 34:7409–7420, 2021.
- Sham M Kakade and Ambuj Tewari. On the generalization ability of online strongly convex programming algorithms. *Advances in Neural Information Processing Systems*, 21, 2008.
- 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.
- Guillaume Lecué and Shahar Mendelson. Aggregation via empirical risk minimization. *Probability Theory and Related Fields*, 145(3-4):591–613, 2009.
- Guillaume Lecué and Philippe Rigollet. Optimal learning with Q-aggregation. *The Annals of Statistics*, 42(1):211–224, 2014.
- Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- Theodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 114–122, 2018.
- Shahar Mendelson. An unrestricted learning procedure. *Journal of the ACM (JACM)*, 66(6):1–42, 2019.

- Carl D. Meyer. *Matrix Analysis and Applied Linear Algebra*. SIAM, 2001.
- Zakaria Mhammedi and Robert C Williamson. Constant regret, generalized mixability, and mirror descent. *Advances in Neural Information Processing Systems*, 31, 2018.
- 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.
- Jaouad Mourtada, Tomas Vaskevicius, and Nikita Zhivotovskiy. Distribution-free robust linear regression. *Mathematical Statistics and Learning*, 4(3-4):253–292, 2021.
- Arkadi Nemirovski. Topics in non-parametric statistics. *Ecole d’Eté de Probabilités de Saint-Flour*, 28:85, 2000.
- Gergely Neu and Nikita Zhivotovskiy. Fast rates for online prediction with abstention. In *Conference on Learning Theory*, pages 3030–3048, 2020.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Francesco Orabona and Nicolò Cesa-Bianchi. Better algorithms for selective sampling. In *International Conference on Machine Learning*, page 433–440, 2011.
- Nikita Puchkin and Nikita Zhivotovskiy. Exponential savings in agnostic active learning through abstention. *IEEE Transactions on Information Theory*, 68(7):4651–4665, 2022.
- Alexander Rakhlin and Karthik Sridharan. Online non-parametric regression. In *Proceedings of The 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pages 1232–1264, 2014.
- Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *Proceedings of the 29th International Conference on Machine Learning*, page 1571–1578, 2012.
- 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.
- El Mehdi Saad and Gilles Blanchard. Fast rates for prediction with limited expert advice. *Advances in Neural Information Processing Systems*, 34:23582–23591, 2021.
- Alexandre B Tsybakov. Optimal rates of aggregation. In *Learning Theory and Kernel Machines*, pages 303–313. Springer, 2003.
- Volodimir G Vovk. Aggregating strategies. In *Proceedings of the 3rd Annual Workshop on Computational Learning Theory*, 1990.
- Volodimir G Vovk. Competitive on-line statistics. *International Statistical Review*, 69(2):213–248, 2001.
- Guanghui Wang, Shiyin Lu, and Lijun Zhang. Adaptivity and optimality: A universal algorithm for online convex optimization. In *Uncertainty in Artificial Intelligence*, pages 659–668, 2020.
- Olivier Wintenberger. Optimal learning with Bernstein online aggregation. *Machine Learning*, 106(1):119–141, 2017.
- Lijun Zhang, Tianbao Yang, and Zhi-Hua Zhou. Dynamic regret of strongly adaptive methods. In *International Conference on Machine Learning*, pages 5882–5891, 2018.
- 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.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]** All the results claimed in the abstract have proofs either in the main body or appendix
  - (b) Did you describe the limitations of your work? **[Yes]** For each result we state the parameters under which the result holds.
  - (c) Did you discuss any potential negative societal impacts of your work? **[No]** Our work is primarily theoretical and we do not foresee any potential negative societal impacts of our work.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]**
  - (b) Did you include complete proofs of all theoretical results? **[Yes]** Although some proofs may only appear in the appendix.
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[N/A]**
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[N/A]**
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[N/A]**
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[N/A]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? **[N/A]**
  - (b) Did you mention the license of the assets? **[N/A]**
  - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[N/A]**
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

## A Details of Section 3 (Our Algorithms)

We restate Lemma 2, after which we prove its result.

**Lemma 2.** For all  $g_1, \dots, g_T \in [-M, M]$  and  $\kappa_0 = \kappa_1 \geq \kappa_2 \cdots \geq \kappa_T$  such that  $\kappa_t \in (0, 1]$ , the predictions  $\hat{y}_t$  of Algorithm 1 run with input  $M$  and  $\eta \in (0, \frac{1}{2}]$  satisfy

$$\sum_{t=1}^T (\hat{y}_t - y_t(i))g_t \leq \frac{M^2 \log(K)}{\kappa_T \eta} + \eta \sum_{t=1}^T \kappa_{t-1} (\hat{y}_t - y_t^*)^2$$

provided  $\max_i |y_t(i) - y_t^*| \leq M$  for all  $t \geq 1$ .

*Proof.* We start by observing that the vector of weights  $\mathbf{p}_t$  is obtained by running lazy EWA with learning rate  $\kappa_t$  on the (signed) surrogate losses  $\tilde{\ell}_t(i) = \gamma(y_t(i) - \hat{y}_t)g_t + \kappa_t(\gamma(y_t(i) - \hat{y}_t)g_t)^2$ . We use  $\kappa_0 = \kappa_1$ . Thus, by [Van der Hoeven et al., 2018, Lemma 1] we have that, for any  $i^* \in [K]$ ,

$$\begin{aligned} & \sum_{t=1}^T (\mathbb{E}_{i \sim \mathbf{p}_t} [\tilde{\ell}_t(i)] - \tilde{\ell}_t(i^*)) \\ & \leq \frac{\log(K)}{\kappa_T} + \sum_{t=1}^T \left( \mathbb{E}_{i \sim \mathbf{p}_t} [\tilde{\ell}_t(i)] + \frac{1}{\kappa_{t-1}} \log \mathbb{E}_{i \sim \mathbf{p}_t} \left[ \exp(-\kappa_{t-1} \tilde{\ell}_t(i)) \right] \right). \end{aligned} \quad (7)$$

Now, using that  $\exp(x - x^2) \leq 1 + x$  for  $|x| \leq \frac{1}{2}$  (Lemma 2.4 by Cesa-Bianchi and Lugosi [2006]) and the fact that  $|\gamma(y_t(i) - \hat{y}_t)g_t| \leq \frac{1}{2}$  due to our choice of  $\gamma$ , we find that

$$\begin{aligned} & \sum_{t=1}^T (\mathbb{E}_{i \sim \mathbf{p}_t} [\tilde{\ell}_t(i)] - \tilde{\ell}_t(i^*)) \\ & \leq \log(K) + \sum_{t=1}^T \left( \mathbb{E}_{i \sim \mathbf{p}_t} [\tilde{\ell}_t(i)] + \frac{1}{\kappa_{t-1}} \log \mathbb{E}_{i \sim \mathbf{p}_t} [1 + \gamma \kappa_{t-1} (\hat{y}_t - y_t(i))g_t] \right) \\ & = \log(K) + \sum_{t=1}^T \mathbb{E}_{i \sim \mathbf{p}_t} [\tilde{\ell}_t(i)], \end{aligned}$$

where the equality is due the fact that since  $\hat{y}_t = \mathbb{E}_{i \sim \mathbf{p}_t} [y_t(i)]$ , we have that  $\mathbb{E}_{i \sim \mathbf{p}_t} [\gamma(\hat{y}_t - y_t(i))g_t] = 0$ . This also implies that  $\mathbb{E}_{i \sim \mathbf{p}_t} [\tilde{\ell}_t(i)] = \mathbb{E}_{i \sim \mathbf{p}_t} [\kappa_{t-1}(\gamma(\hat{y}_t - y_t(i))g_t)^2]$ . Thus, we may write

$$\begin{aligned} & \mathbb{E}_{i \sim \mathbf{p}_t} [\tilde{\ell}_t(i)] - \tilde{\ell}_t(i^*) \\ & = \mathbb{E}_{i \sim \mathbf{p}_t} [\kappa_{t-1}(\gamma(\hat{y}_t - y_t(i))g_t)^2] + \gamma(\hat{y}_t - y_t(i))g_t - \kappa_{t-1}(\gamma(\hat{y}_t - y_t(i))g_t)^2. \end{aligned}$$

Combining with the above and reordering we find

$$\sum_{t=1}^T \gamma(\hat{y}_t - y_t(i))g_t \leq \frac{\log(K)}{\kappa_T} + \sum_{t=1}^T \kappa_{t-1} (\gamma(\hat{y}_t - y_t(i))g_t)^2 \quad (8)$$

After dividing both sides by  $\gamma = \frac{\eta}{M^2}$ , this gives us

$$\begin{aligned} \sum_{t=1}^T (\hat{y}_t - y_t(i))g_t & \leq \frac{M^2 \log(K)}{\kappa_T \eta} + \frac{\eta}{M^2} \sum_{t=1}^T \kappa_{t-1} ((\hat{y}_t - y_t(i))g_t)^2 \\ & \leq \frac{M^2 \log(K)}{\kappa_T \eta} + \eta \sum_{t=1}^T \kappa_{t-1} (\hat{y}_t - y_t(i^*))^2. \end{aligned}$$

completing the proof.  $\square$

We now restate Theorem 1 and provide its proof.

**Theorem 1.** Consider the squared loss and two experts  $y_t(1) = 0$  and  $y_t(2) = 1$  for all  $t \geq 1$ . Let  $\hat{y}_t^{\text{EWA}}$  be the EWA predictions. There is a sequence  $y_1, y_2, \dots$  such that  $y_t \in [0, 1]$ ,  $t \geq 1$  and, for large enough  $T$ , the regret of EWA with  $\eta = \frac{1}{2}$  satisfies  $-12 \log T \leq \mathcal{R}_T \leq 2 \log 2$  and, at the same time,

$$\sum_{t=1}^T (\hat{y}_t^{\text{EWA}} - y_t^*)^2 \geq T/2.$$

*Proof.* The proof uses a construction similar to one used in Audibert [2007]. Our idea is to show that for some environments the output of EWA is close to the *follow the leader* prediction. This can lead to large variance when we follow a wrong expert for most of the rounds. For  $T$  large enough consider the following sequence:

$$y_t = \begin{cases} 3/4, & \text{if } t \leq 4\lceil \log T \rceil, \\ 1/2, & \text{if } 4\lceil \log T \rceil < t < T - 8\lceil \log T \rceil, \\ 1/4, & \text{if } T - 8\lceil \log T \rceil \leq t \leq T. \end{cases}$$

Fix  $\eta = 1/2$ . Since  $y_t, y_t(1), y_t(2) \in [0, 1]$ , and the squared loss is  $1/2$ -exp-concave on this domain (see Appendix E) we have  $\mathcal{R}_T \leq 2 \log 2$ .

Next, we show the lower bound. Since  $y_t = 1/4$  appears more frequently in the sequence, we have that  $y_t(1) = 0$  is the prediction of the best expert. However, until the last  $8\lceil \log T \rceil$  rounds, that is, for any  $4\lceil \log T \rceil < t' < T - 8\lceil \log T \rceil$  the EWA algorithm puts most of its weight on the second expert predicting  $y_t(2) = 1$ . At the same time, both experts suffer the same loss when  $y_t = 1/2$ . Formally, for any such  $t'$  we have

$$\sum_{t=1}^{t'} (y_t - y_t(1))^2 - \sum_{t=1}^{t'} (y_t - y_t(2))^2 = 4\lceil \log T \rceil (9/16 - 1/16) = 2\lceil \log T \rceil.$$

Therefore, for the same  $t'$ , the weight of the first expert in the EWA prediction with  $\eta = 1/2$  is

$$p_{t'}(1) = \frac{\exp\left(-\sum_{t=1}^{t'} (y_t - y_t(1))^2/2\right)}{\sum_{i=1}^2 \exp\left(-\sum_{t=1}^{t'} (y_t - y_t(i))^2/2\right)} = \frac{1}{1 + \exp(\lceil \log T \rceil)} \leq \frac{1}{1 + T}.$$

Thus, we have

$$\hat{y}_{t'}^{\text{EWA}} = p_{t'}(1)y_{t'}(1) + p_{t'}(2)y_{t'}(2) \geq \frac{T}{T+1}.$$

We are ready to bound the regret. Our idea will be just to use the boundedness of the loss when  $y_t \in \{1/4, 3/4\}$  and compute the regret over remaining rounds. Using elementary algebra, we have

$$\begin{aligned} \mathcal{R}_T &\geq (T - 12\lceil \log T \rceil) \cdot \left( \left( \frac{T}{T+1} - 1/2 \right)^2 - (1/2)^2 \right) - 12\lceil \log T \rceil (3/4)^2 \\ &\geq \frac{-T(T - 12\lceil \log T \rceil)}{(T+1)^2} - 12(3/4)^2 \lceil \log T \rceil \geq -12 \log T, \end{aligned}$$

for all  $T \geq 4$ . At the same time, the following variance bound holds

$$\sum_{t=1}^T (\hat{y}_t^{\text{EWA}} - y_t(1))^2 \geq \sum_{t=4\lceil \log T \rceil+1}^{T-8\lceil \log T \rceil-1} (\hat{y}_t^{\text{EWA}} - y_t(1))^2 \geq (T - 12\lceil \log T \rceil - 2) \left( \frac{T}{T+1} \right)^2 \geq \frac{T}{2},$$

provided that  $T > 150$ . The claim follows.  $\square$

Here we restate Lemma 3, after which we prove it.

**Lemma 3.** For all  $g_1, \dots, g_T \in \mathbb{R}$  and  $\kappa_1 \geq \dots \geq \kappa_T \in (0, 1]$ , the predictions  $\hat{y}_t$  of Algorithm 2 run with inputs  $\eta > 0$ ,  $\sigma = D^2$ ,  $G \geq \max_t |g_t|$ , and  $Z > 0$

$$\sum_{t=1}^T (\hat{y}_t - y_t(\mathbf{u}))g_t \leq \frac{dG^2}{2\kappa_T\eta} \log \left( 1 + D^2\eta^2 \left( \max_{t=1, \dots, T} \|\mathbf{x}_t\|_2^2 \right) \frac{T}{d} \right) + \frac{G^2}{2\eta} + \eta \sum_{t=1}^T \kappa_t (\hat{y}_t - y_t(\mathbf{u}))^2,$$

for any  $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{R}^d$ , and for any  $\mathbf{u} \in \mathcal{W}_T \equiv \bigcap_{t=1}^T \{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x}_t \rangle| \leq Z\}$  such that  $\|\mathbf{u}\|_2 \leq D$ .

*Proof.* We start by observing that  $\mathbf{w}_t$  is the mean of continuous exponential weights with a Gaussian prior and learning rate 1 on (signed) surrogate losses  $\tilde{\ell}_t^{\text{or}}(\mathbf{w}) = \langle \mathbf{w} - \mathbf{w}_t, \mathbf{z}_t \rangle + \kappa_t (\mathbf{w} - \mathbf{w}_t)^\top \mathbf{z}_t \mathbf{z}_t^\top (\mathbf{w} - \mathbf{w}_t)$ , see [Van der Hoeven et al., 2018, Section 4]. Thus, for any  $\mathbf{u} \in \mathcal{W}_T$ , by Van der Hoeven et al. [2018, Theorem 5] we have that

$$\sum_{t=1}^T (\tilde{\ell}_t^{\text{or}}(\mathbf{w}_t) - \tilde{\ell}_t^{\text{or}}(\mathbf{u})) \leq \frac{\|\mathbf{u}\|_2^2}{2\sigma} + \frac{1}{2} \sum_{t=1}^T \mathbf{z}_t^\top \Sigma_{t+1} \mathbf{z}_t.$$

Using that  $\kappa_1, \dots, \kappa_T \in (0, 1]$  and the Sherman-Morrison formula to compute the inverse we find that

$$\begin{aligned} 2\mathbf{z}_t^\top \left( 2 \sum_{s=1}^t \kappa_s \mathbf{z}_s \mathbf{z}_s^\top + \frac{1}{\sigma} I \right)^{-1} \mathbf{z}_t &\leq \frac{1}{\kappa_T} 2\mathbf{z}_t^\top \left( 2 \sum_{s=1}^t \mathbf{z}_s \mathbf{z}_s^\top + \frac{1}{\sigma} I \right)^{-1} \mathbf{z}_t \\ &= \frac{1}{\kappa_T} \left( 2\mathbf{z}_t^\top \Sigma_t \mathbf{z}_t - \frac{(2\mathbf{z}_t^\top \Sigma_t \mathbf{z}_t)^2}{1 + 2\mathbf{z}_t^\top \Sigma_t \mathbf{z}_t} \right) \\ &= \frac{1}{\kappa_T} \left( 1 - \frac{1}{1 + 2\mathbf{z}_t^\top \Sigma_t \mathbf{z}_t} \right) \\ &\leq \frac{1}{\kappa_T} \log(1 + 2\mathbf{z}_t^\top \Sigma_t \mathbf{z}_t) = \frac{1}{\kappa_T} \log \left( \frac{\text{Det}(\Sigma_{t+1})}{\text{Det}(\Sigma_t)} \right), \end{aligned}$$

where the second inequality is due to the fact that  $1 - 1/x \leq \log(x)$  for  $x > 0$  and the final equality can be found on, for example, in Meyer [2001, page 475]. Thus, we have that

$$\begin{aligned} \sum_{t=1}^T \mathbf{z}_t^\top \Sigma_{t+1} \mathbf{z}_t &\leq \frac{1}{\kappa_T} \sum_{t=1}^T \frac{1}{2} \log \left( \frac{\text{Det}(\Sigma_{t+1})}{\text{Det}(\Sigma_t)} \right) \\ &\leq \frac{1}{2\kappa_T} \log \left( \frac{\text{Det}(\Sigma_{T+1})}{\text{Det}(\Sigma_1)} \right) = \frac{1}{2\kappa_T} \log \text{Det} \left( I + \sigma \sum_{t=1}^T 2\mathbf{z}_t \mathbf{z}_t^\top \right). \end{aligned}$$

Now, following the proof and discussion of Cesa-Bianchi and Lugosi [2006, Theorem 11.2] we have that

$$\begin{aligned} \log \text{Det} \left( I + \sigma^2 \sum_{t=1}^T 2\mathbf{z}_t \mathbf{z}_t^\top \right) &\leq d \log \left( 1 + 2\sigma \left( \max_t \|\mathbf{z}_t\|_2^2 \right) \frac{T}{d} \right) \\ &\leq d \log \left( 1 + 2\sigma\gamma^2 \left( \max_t \|\mathbf{x}_t g_t\|_2^2 \right) \frac{T}{d} \right). \end{aligned}$$

By combining the above we find

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t g_t \rangle &= \frac{1}{\gamma} \sum_{t=1}^T (\langle \mathbf{w}_t - \mathbf{u}, \mathbf{z}_t \rangle - \kappa_t \langle \mathbf{w}_t - \mathbf{u}, \mathbf{z}_t \rangle^2) + \sum_{t=1}^T \gamma \kappa_t (\langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t g_t \rangle)^2 \\ &\leq \frac{\|\mathbf{u}\|_2^2}{2\sigma\gamma} + \frac{d}{2\gamma\kappa_T} \log \left( 1 + 2\sigma\gamma^2 \left( \max_t \|\mathbf{x}_t g_t\|_2^2 \right) \frac{T}{d} \right) + \sum_{t=1}^T \gamma \kappa_t (\langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t g_t \rangle)^2. \end{aligned} \tag{9}$$



Now, we continue by using that  $\langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t \rangle = \hat{y}_t - y_t(\mathbf{u})$ ,  $\gamma = \frac{\eta}{G^2}$ ,  $\sigma = D^2$ , and that  $\|\mathbf{u}\|_2 \leq D$

$$\begin{aligned} & \sum_{t=1}^T (\hat{y}_t - y_t(\mathbf{u})) g_t \\ & \leq \frac{1}{2\gamma} + \frac{dG^2}{2\kappa_T\eta} \log \left( 1 + 2D^2\gamma^2 \left( \max_t \|\mathbf{x}_t g_t\|_2^2 \right) \frac{T}{d} \right) + \sum_{t=1}^T \gamma \kappa_t ((\hat{y}_t - y_t(\mathbf{u})) g_t)^2 \\ & \leq \frac{G^2}{2\eta} + \frac{dG^2}{2\kappa_T\eta} \log \left( 1 + 2D^2\eta^2 \left( \max_t \|\mathbf{x}_t\|_2^2 \right) \frac{T}{d} \right) + \eta \sum_{t=1}^T \kappa_t (\hat{y}_t - y_t(\mathbf{u}))^2, \end{aligned}$$

where we used the assumption that  $g_t^2 \leq G^2$ , which completes the proof.  $\square$

## B Details of Section 4 (Statistical Learning)

To directly use our results obtained in the general online setting, we set the following notation for the rest of the section:

$$\ell_t(\cdot) = \ell(\cdot, Y_t), \quad y_t(f) = f(X_t), \quad y_t^* = f^*(X_t) \quad \text{and} \quad \hat{y}_t = \hat{f}_t(X_t), \quad (10)$$

where  $\hat{f}_t$  is a statistical estimator constructed based on  $(X_1, Y_1), \dots, (X_{t-1}, Y_{t-1})$ .

Let  $r_t = \ell_t(\hat{y}_t) - \ell_t(y_t^*)$ . To prove our high-probability bounds, we are interested in controlling  $\sum_{t=1}^T \mathbb{E}_{t-1}[r_t]$ . As we mentioned, the challenge in obtaining high-probability bounds comes from bounding  $\mathbb{E}_{t-1}[r_t] - r_t$ , which may be of order  $\sqrt{T}$  in the worst case due to the variance of  $r_t$ . The negative term in our regret bounds is used to control the variance term.

The following two versions of Freedman's inequality for martingales appear explicitly in [Beygelzimer et al., 2011, Theorem 1] and [Rakhlin et al., 2012, Lemma 3]. We use them to prove Lemmas 5 and 6, which in turn are used to prove Theorems 8, 9, and 10.

**Lemma 4** (Versions of Freedman's inequality). *Let  $X_1, \dots, X_T$  be a martingale difference sequence adapted to a filtration  $(\mathcal{F}_i)_{i \leq T}$ . That is, in particular,  $\mathbb{E}_{t-1}[X_t] = 0$ . Suppose that  $|X_t| \leq R$  almost surely. Then for any  $\delta \in (0, 1)$ ,  $\lambda \in [0, 1/R]$ , with probability at least  $1 - \delta$ , it holds that*

$$\sum_{t=1}^T X_t \leq \lambda(e-2) \sum_{t=1}^T \mathbb{E}_{t-1}[X_t^2] + \frac{\log(1/\delta)}{\lambda}. \quad (11)$$

Moreover, if  $\delta \in (0, 1/2)$ ,  $T \geq 4$ , then uniformly over all  $s \leq T$ , with probability at least  $1 - \delta$ , it holds that

$$\sum_{t=1}^s X_t \leq 4 \sqrt{\sum_{t=1}^s \mathbb{E}_{t-1}[X_t^2] \log(\log(T)/\delta)} + 2R \log(\log(T)/\delta). \quad (12)$$

**Lemma 5.** *Under the notation (10) suppose that  $\max_t \max\{|\hat{y}_t|, |y_t^*|\} \leq \frac{1}{2}M$  almost surely and that  $\max_t |\ell'_t(y)| \leq M$  almost surely for all  $y$  such that  $|y| \leq \frac{1}{2}M$ . Suppose that  $\hat{y}_1, \dots, \hat{y}_T$  satisfy (2) with  $\eta = \frac{\mu}{4}$ . Then, with probability at least  $1 - \delta$ , it holds that*

$$\sum_{t=1}^T \mathbb{E}_{t-1}[\ell_t(\hat{y}_t) - \ell_t(y_t^*)] \leq \frac{8C_T}{\mu} + B_T + \log(1/\delta) \min \left\{ \frac{1}{(2 + \frac{\mu}{2})M^2}, \frac{\mu}{6M^2} \left( 1 + \frac{\mu^2}{16} \right)^{-1} \right\}^{-1}.$$

*Proof.* Let  $v_t = (\hat{y}_t - y_t^*)^2$  and let  $r_t = \ell_t(\hat{y}_t) - \ell_t(y_t^*)$ . By convexity and the assumptions on  $y_t^*$ ,  $\hat{y}_t$ , and  $\ell'_t$  we have that  $|\ell_t(\hat{y}_t) - \ell_t(y_t^*)| \leq |\hat{y}_t - y_t^*| M \leq M^2$ . This implies that

$$\left| r_t - \frac{\mu}{4} v_t \right| \leq |r_t| + \frac{\mu}{4} v_t \leq M^2 + \frac{\mu}{4} M^2.$$

Thus, by equation (11) we have that, for  $\lambda \in \left[0, \frac{1}{2(M^2 + \frac{\mu}{4}M^2)}\right]$ , with probability at least  $1 - \delta$ ,

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{t-1} \left[ r_t + \frac{\mu}{4} v_t \right] - \frac{\mu}{4} v_t - r_t \\ & \leq \lambda(e-2) \sum_{t=1}^T \mathbb{E}_{t-1} \left[ \left( \mathbb{E}_{t-1} \left[ r_t + \frac{\mu}{4} v_t \right] - \frac{\mu}{4} v_t - r_t \right)^2 \right] + \frac{\log(1/\delta)}{\lambda} \\ & \leq \lambda(e-2) \sum_{t=1}^T \mathbb{E}_{t-1} \left[ \left( \frac{\mu}{4} v_t + r_t \right)^2 \right] + \frac{\log(1/\delta)}{\lambda}. \end{aligned}$$

where we used that  $\mathbb{E}[(X - \mathbb{E}[X])^2] \leq \mathbb{E}[X^2]$ . Since  $r_t^2 \leq (y_t^* - \hat{y}_t)^2 M^2$  we have that

$$\begin{aligned} \mathbb{E}_{t-1} \left[ \left( r_t + \frac{\mu}{4} v_t \right)^2 \right] & \leq 2 \mathbb{E}_{t-1}[r_t^2] + \frac{\mu^2}{8} \mathbb{E}[v_t^2] \\ & \leq \mathbb{E}_{t-1} \left[ 2(y_t^* - \hat{y}_t)^2 M^2 + \frac{\mu^2 M^2}{8} (y_t^* - \hat{y}_t)^2 \right] \\ & = 2 \left( M^2 + \frac{\mu^2 M^2}{16} \right) \mathbb{E}_{t-1}[v_t]. \end{aligned}$$

which means that, with probability at least  $1 - \delta$ ,

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{t-1} \left[ r_t + \frac{\mu}{4} v_t \right] - \frac{\mu}{4} v_t - r_t \\ & \leq \lambda(e-2) \sum_{t=1}^T 2 \left( M^2 + \frac{\mu^2 M^2}{16} \right) \mathbb{E}_{t-1}[v_t] + \frac{\log(1/\delta)}{\lambda} \tag{13} \\ & \leq \lambda \sum_{t=1}^T \frac{3}{2} \left( M^2 + \frac{\mu^2 M^2}{16} \right) \mathbb{E}_{t-1}[v_t] + \frac{\log(1/\delta)}{\lambda} \end{aligned}$$

Using the guarantee on  $\mathcal{R}_T$  in equation (2), Lemma 1, and replacing  $\eta$  with  $\frac{\mu}{4}$  we have that

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{t-1}[r_t] \\ & = \mathcal{R}_T + \sum_{t=1}^T (\mathbb{E}_{t-1}[r_t] - r_t) \\ & \leq \frac{C_T}{\eta} + \left( \eta - \frac{\mu}{2} \right) \sum_{t=1}^T v_t + B_T + \sum_{t=1}^T (\mathbb{E}_{t-1}[r_t] - r_t) \\ & = \frac{C_T}{\eta} + \left( \eta - \frac{\mu}{4} \right) \sum_{t=1}^T v_t + B_T - \frac{\mu}{4} \sum_{t=1}^T \mathbb{E}_{t-1}[v_t] + \sum_{t=1}^T \left( \mathbb{E}_{t-1}[r_t] - r_t + \frac{\mu}{4} (\mathbb{E}_{t-1}[v_t] - v_t) \right) \\ & = \frac{8C_T}{\mu} + B_T - \frac{\mu}{4} \sum_{t=1}^T \mathbb{E}_{t-1}[v_t] + \sum_{t=1}^T \left( \mathbb{E}_{t-1}[r_t] - r_t + \frac{\mu}{4} (\mathbb{E}_{t-1}[v_t] - v_t) \right). \end{aligned}$$

Thus, by equation (13) we have that, with probability at least  $1 - \delta$ ,

$$\sum_{t=1}^T \mathbb{E}_{t-1}[r_t] \leq \frac{8C_T}{\mu} + B_T - \frac{\mu}{4} \sum_{t=1}^T \mathbb{E}_{t-1}[v_t] + \lambda \sum_{t=1}^T \frac{3}{2} \left( M^2 + \frac{\mu^2 M^2}{16} \right) \mathbb{E}_{t-1}[v_t] + \frac{\log(1/\delta)}{\lambda}.$$

Thus, setting  $\lambda = \min \left\{ \frac{1}{2(M^2 + \frac{\mu}{4}M^2)}, \frac{\mu}{6} \left( M^2 + \frac{\mu^2 M^2}{16} \right)^{-1} \right\}$  completes the proof.  $\square$

**Lemma 6.** *Under the notation (10) suppose that  $\max_t \max\{|\hat{y}_t|, |y_t^*|\} \leq \frac{1}{2}M$  almost surely and that  $\max_t |\ell'_t(y)| \leq M$  almost surely for all  $y$  such that  $|y| \leq \frac{1}{2}M$ . Suppose that  $\hat{y}_1, \dots, \hat{y}_S$  satisfy (2) with  $\eta = \frac{\mu}{8}$ . Then, for  $\delta \in (0, \frac{1}{2})$ ,  $T \geq 4$ , and uniformly over all  $S \leq T$ , with probability at least  $1 - \delta$ , it holds that*

$$\sum_{t=1}^S \mathbb{E}_{t-1} [\ell_t(\hat{y}_t) - \ell_t(y_t^*)] \leq \frac{8C_T}{\mu} + B_T - \frac{\mu}{8} \sum_{t=1}^S (\hat{y}_t - y_t^*)^2 + M^2 \left( \frac{32}{\mu} + 3\mu + 4 \right) \log \left( \frac{\log T}{\delta} \right).$$

*Proof.* Let  $v_t = (\hat{y}_t - y_t^*)^2$  and  $r_t = \ell_t(\hat{y}_t) - \ell_t(y_t^*)$ . By convexity and the assumptions on  $y_t^*, \hat{y}_t$ , and  $\ell'_t$  we have that  $|\ell_t(\hat{y}_t) - \ell_t(y_t^*)| \leq |\hat{y}_t - y_t^*| M \leq M^2$ . This implies that

$$\left| r_t - \frac{\mu}{4} v_t \right| \leq |r_t| + \frac{\mu}{4} v_t \leq M^2 + \frac{\mu}{4} M^2.$$

Thus, by equation (12) we have that, with probability at least  $1 - \delta$

$$\begin{aligned} & \sum_{t=1}^S \mathbb{E}_{t-1} \left[ r_t + \frac{\mu}{4} v_t \right] - \frac{\mu}{4} v_t - r_t \\ & \leq 4 \sqrt{\sum_{t=1}^S \mathbb{E}_{t-1} \left[ \left( \frac{\mu}{4} v_t + r_t \right)^2 \right]} \log(\log(T)/\delta) + M^2 (4 + \mu) \log(\log(T)/\delta). \end{aligned}$$

where we used that  $\mathbb{E}[(X - \mathbb{E}[X])^2] \leq \mathbb{E}[X^2]$ . Since  $r_t^2 \leq (y_t^* - \hat{y}_t)^2 M^2$  we have that

$$\begin{aligned} \mathbb{E}_{t-1} \left[ \left( r_t + \frac{\mu}{4} v_t \right)^2 \right] & \leq 2 \mathbb{E}_{t-1} [r_t^2] + \frac{2\mu^2}{16} \mathbb{E}_{t-1} [v_t^2] \\ & \leq \mathbb{E}_{t-1} \left[ 2(y_t^* - \hat{y}_t)^2 M^2 + \frac{2\mu^2 M^2}{16} (y_t^* - \hat{y}_t)^2 \right] = 2 \left( M^2 + \frac{\mu^2 M^2}{16} \right) \mathbb{E}_{t-1} [v_t]. \end{aligned}$$

which means that with probability at least  $1 - \delta$

$$\begin{aligned} & \sum_{t=1}^S \mathbb{E}_{t-1} \left[ r_t + \frac{\mu}{4} v_t \right] - \frac{\mu}{4} v_t - r_t \\ & \leq 4 \sqrt{\sum_{t=1}^S 2 \left( M^2 + \frac{\mu^2 M^2}{16} \right) \mathbb{E}_{t-1} [v_t] \log(\log(T)/\delta) + 2R \log(\log(T)/\delta)} \quad (14) \\ & \leq \lambda \frac{1}{2} \sum_{t=1}^S \mathbb{E}_{t-1} [v_t] + \frac{16}{\lambda} \left( M^2 + \frac{\mu^2 M^2}{16} \right) \log(\log(T)/\delta) + M^2 (4 + \mu) \log(\log(T)/\delta) \end{aligned}$$

for any  $\lambda > 0$ , where in the final inequality we used  $\sqrt{ab} = \frac{1}{2} \inf_{\eta > 0} \eta a + \frac{b}{\eta}$  for  $a, b \geq 0$ . Using the guarantee on  $\mathcal{R}_T$  in equation (2), Lemma 1, and replacing  $\eta$  with  $\frac{\mu}{8}$  we have that

$$\begin{aligned} \sum_{t=1}^S \mathbb{E}_{t-1} [r_t] & = \mathcal{R}_T + \sum_{t=1}^S (\mathbb{E}_{t-1} [r_t] - r_t) \\ & \leq \frac{C_T}{\eta} + \left( \eta - \frac{\mu}{2} \right) \sum_{t=1}^S v_t + B_T + \sum_{t=1}^S (\mathbb{E}_{t-1} [r_t] - r_t) \\ & = \frac{C_T}{\eta} + \left( \eta - \frac{\mu}{4} \right) \sum_{t=1}^S v_t + B_T - \frac{\mu}{4} \sum_{t=1}^S \mathbb{E}_{t-1} [v_t] \\ & \quad + \sum_{t=1}^S \left( \mathbb{E}_{t-1} [r_t] - r_t + \frac{\mu}{4} (\mathbb{E}_{t-1} [v_t] - v_t) \right) \\ & = \frac{8C_T}{\mu} + B_T - \frac{\mu}{8} \sum_{t=1}^S v_t - \frac{\mu}{4} \sum_{t=1}^S \mathbb{E}_{t-1} [v_t] \\ & \quad + \sum_{t=1}^S \left( \mathbb{E}_{t-1} [r_t] - r_t + \frac{\mu}{4} (\mathbb{E}_{t-1} [v_t] - v_t) \right). \end{aligned}$$

Thus, by (14) we have that, with probability at least  $1 - \delta$ ,

$$\begin{aligned} \sum_{t=1}^S \mathbb{E}_{t-1}[r_t] &\leq \frac{8C_T}{\mu} + B_T - \frac{\mu}{8} \sum_{t=1}^S v_t - \frac{\mu}{4} \sum_{t=1}^S \mathbb{E}_{t-1}[v_t] + \lambda \frac{1}{2} \sum_{t=1}^S \mathbb{E}_{t-1}[v_t] \\ &\quad + \frac{16}{\lambda} \left( M^2 + \frac{\mu^2 M^2}{16} \right) \log(\log(T)/\delta) + M^2 (4 + \mu) \log(\log(T)/\delta) \end{aligned}$$

Thus, setting  $\lambda = \frac{\mu}{2}$  gives us

$$\sum_{t=1}^S \mathbb{E}_{t-1}[r_t] \leq \frac{8C_T}{\mu} + B_T - \frac{\mu}{8} \sum_{t=1}^S v_t + M^2 \left( \frac{32}{\mu} + 3\mu + 4 \right) \log(\log(T)/\delta),$$

which completes the proof.  $\square$

The following Theorem is the detailed statement of Theorem 2 in the main body of the paper.

**Theorem 8.** *Suppose that for all  $f \in \mathcal{F}$   $|f(X)| \leq \frac{1}{2}M$  almost surely, that  $|\partial_y \ell(y, Y)| \leq M$  almost surely for all  $y$  such that  $|y| \leq \frac{1}{2}M$ , and that  $\ell$  is  $\mu$ -strongly convex in its first argument. Then, with probability at least  $1 - \delta$ , Algorithm 3 with input parameters  $T$ ,  $S = \frac{8M^2 \log(K)}{\mu} + M^2 \left( \frac{32}{\mu} + 3\mu + 4 \right) \log(\log(T)/\delta)$ ,  $\eta = \frac{\mu}{8}$ , and  $M$  guarantees*

$$R(\hat{f}) - \min_{f \in \mathcal{F}} R(f) \leq \begin{cases} 0 & \text{if } S < T, \\ \frac{M^2}{\mu T} \left( 8 \log(K) + (32 + 3\mu^2 + 4\mu) \log(\log(T)/\delta) \right) & \text{if } S = T. \end{cases}$$

*Proof.* Convexity of  $R$  gives us

$$R(\hat{f}) \leq \frac{1}{S} \sum_{t=1}^S R \left( \sum_{i=1}^K p_t(i) f_i \right) = \frac{1}{S} \sum_{t=1}^S \mathbb{E}_{t-1} \left[ \ell_t \left( \sum_{i=1}^K p_t(i) f_i(X_t) \right) \right] = \frac{1}{S} \sum_{t=1}^S \mathbb{E}_{t-1}[\ell_t(\hat{y}_t)].$$

Now, for any fixed  $f \in \mathcal{F}$ , by Lemma 6 and Lemma 2 we have that, with probability  $1 - \delta$ , simultaneously for all  $S \leq T$ ,

$$\begin{aligned} R(\hat{f}) - R(f) &\leq \frac{1}{S} \sum_{t=1}^S \mathbb{E}_{t-1}[\ell_t(\hat{y}_t) - \ell_t(y_t(i))] \\ &\leq \frac{\frac{8M^2 \log(K)}{\mu} - \frac{\mu}{8} \sum_{t=1}^S (\hat{y}_t - f(X_t))^2 + M^2 \left( \frac{32}{\mu} + 3\mu + 4 \right) \log(\log(T)/\delta)}{S}. \end{aligned} \tag{15}$$

We split the remainder of the proof into two cases. Either  $S < T$  or  $S = T$ . If  $S < T$  then  $S = \frac{8M^2 \log(K)}{\mu} + M^2 \left( \frac{32}{\mu} + 3\mu + 4 \right) \log(\log(T)/\delta) \leq \frac{\mu}{8} \min_i \sum_{t=1}^S (\hat{y}_t - y_t(i))$  and thus

$$R(\hat{f}) - R(f) \leq 0.$$

To complete the proof, observe that if  $S = T$ , then from (15) we have that

$$R(\hat{f}) - R(f) \leq \frac{\frac{8M^2 \log(K)}{\mu} + M^2 \left( \frac{32}{\mu} + 3\mu + 4 \right) \log(\log(T)/\delta)}{T}.$$

The claim follows.  $\square$

The following result is a detailed version of Theorem 3 in the main text.

**Theorem 9.** *Suppose that for all  $f \in \mathcal{F}$   $|f(X)| \leq \frac{1}{2}M$  almost surely, that  $|\partial_y \ell(y, Y)| \leq M$  almost surely for all  $y$  s.t.  $|y| \leq \frac{1}{2}M$ , and that  $\ell$  is  $\mu$ -strongly convex in its first argument. Then, with probability at least  $1 - \delta$ , Algorithm 3 with input parameters  $T$ ,  $\eta = \frac{\mu}{4}$ ,  $S = \infty$ , and  $M$ ,*

$$\begin{aligned} &R(\hat{f}) - \min_{f \in \mathcal{F}} R(f) \\ &\leq \frac{1}{T} \left( \frac{4M^2 \log(K)}{\mu} + \log(1/\delta) \min \left\{ \frac{1}{\left(2 + \frac{\mu}{2}\right)M^2}, \frac{\mu}{6M^2} \left(1 + \frac{\mu^2}{16}\right)^{-1} \right\}^{-1} \right). \end{aligned}$$

*Proof.* Observe that  $\eta < \frac{1}{2}$  by assumption on  $\mu$ , making it a valid choice. Furthermore, by the choice of  $S$  we have that on termination of Algorithm 3  $S = T$  and thus convexity of  $R$  gives us

$$R(\hat{f}) \leq \frac{1}{T} \sum_{t=1}^T R \left( \sum_{i=1}^K p_t(i) f_i \right) = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{t-1} \left[ \ell_t \left( \sum_{i=1}^K p_t(i) f_i(X_t) \right) \right] = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{t-1} [\ell_t(\hat{y}_t)].$$

Now, for any  $f \in \mathcal{F}$ , by Lemma 5 and Lemma 2 we have that with probability  $1 - \delta$

$$\begin{aligned} R(\hat{f}) - R(f) &\leq \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{t-1} [\ell_t(\hat{y}_t) - \ell_t(y_t(i))] \\ &\leq \frac{1}{T} \left( \frac{4M^2 \log(K)}{\mu} + \log(1/\delta) \min \left\{ \frac{1}{(2 + \frac{\mu}{2})M^2}, \frac{\mu}{6M^2} \left( 1 + \frac{\mu^2}{16} \right)^{-1} \right\}^{-1} \right), \end{aligned}$$

which completes the proof.  $\square$

The following Theorem is a detailed version of Theorem 4.

**Theorem 10.** Fix  $Z > 0$ . Suppose that  $\ell$  is  $\mu$ -strongly convex in its first argument and that  $\|X_t\|_2 \leq B$  and  $\sup_{y \in [-Z, Z]} |\partial \ell(y, T)| \leq G$  almost surely. For any  $\mathbf{w} \in \{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x} \rangle| \leq Z \text{ for all } \mathbf{x} \in \mathcal{X}\}$  such that  $\|\mathbf{w}\|_2 \leq D$ , with probability at least  $1 - \delta$

$$\begin{aligned} R \left( \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t \right) - R(\mathbf{w}) &\leq \frac{1}{T} \left( \frac{8G^2 + \frac{dG^2}{2} \log \left( 1 + D^2 \mu^2 B^2 \frac{T}{2d} \right)}{\mu} \right. \\ &\quad \left. + \log(1/\delta) \min \left\{ \frac{1}{(2 + \frac{\mu}{2}) \max\{G, 2Z\}^2}, \frac{\mu}{6 \max\{G, 2Z\}^2} \left( 1 + \frac{\mu^2}{16} \right)^{-1} \right\}^{-1} \right), \end{aligned}$$

where  $\mathbf{w}_t$  are given by Algorithm 2 with  $\eta = \frac{\mu}{4}$ ,  $\sigma = D^2$ ,  $G$ ,  $\kappa_t = 1$ , and feedback  $g_t = \ell'_t(\langle \mathbf{w}_t, X_t \rangle)$  for  $t = 1, \dots, T$ .

*Proof.* Convexity of  $R$  together with  $\mathbb{E}_{t-1}[\ell_t(\hat{y}_t)] = R(\mathbf{w}_t)$  gives us

$$R(\bar{\mathbf{w}}) \leq \frac{1}{T} \sum_{t=1}^T R(\mathbf{w}_t) = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{t-1}[\ell_t(\hat{y}_t)].$$

Let  $y_t(\mathbf{w}) = \langle \mathbf{w}, X_t \rangle$ . Let  $M = \max\{G, 2Z\}$ . Since  $\{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x} \rangle| \leq Z \text{ for all } \mathbf{x} \in \mathcal{X}\} \subseteq \mathcal{W}_T$ , by Lemma 5 and Lemma 3, we have that for any fixed  $\mathbf{w} \in \{\mathbf{w} : \|\mathbf{w}\|_2 \leq D \text{ and } |\langle \mathbf{w}, X \rangle| \leq Z \text{ for all } \mathbf{x} \in \mathcal{X}\}$ , with probability at least  $1 - \delta$ ,

$$\begin{aligned} R(\bar{\mathbf{w}}) - R(\mathbf{w}) &\leq \frac{1}{T} \sum_{t=1}^T (\mathbb{E}_{t-1}[\ell_t(\hat{y}_t)] - \mathbb{E}_{t-1}[\ell_t(y_t(\mathbf{w}))]) \\ &\leq \frac{1}{T} \left( \frac{8C_T}{\mu} + \log(1/\delta) \min \left\{ \frac{1}{(2 + \frac{\mu}{2})M^2}, \frac{\mu}{6M^2} \left( 1 + \frac{\mu^2}{16} \right)^{-1} \right\}^{-1} \right) \\ &= \frac{1}{T} \left( \frac{8G^2 + \frac{dG^2}{2} \log \left( 1 + D^2 \mu^2 B^2 \frac{T}{2d} \right)}{\mu} \right. \\ &\quad \left. + \log(1/\delta) \min \left\{ \frac{1}{(2 + \frac{\mu}{2})M^2}, \frac{\mu}{6M^2} \left( 1 + \frac{\mu^2}{16} \right)^{-1} \right\}^{-1} \right). \end{aligned}$$

$\square$

## C Details of Section 5 (Corrupted Feedback)

We first restate Theorem 5, after which we prove it.

**Theorem 5.** Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses and corruptions  $c_1, \dots, c_T \in \mathbb{R}$ . Then the predictions  $\hat{y}_t$  of Algorithm 1 run with inputs  $M \geq \max_t |\ell'_t(\hat{y}_t) + c_t|$ ,  $\eta = \frac{\mu}{4}$ , feedback  $g_t = \ell'_t(\hat{y}_t) + c_t$ , and  $\kappa_t = 1$  satisfy

$$\mathcal{R}_T \leq \frac{8M^2 \log(K)}{\mu} + \sum_{t=1}^T \frac{c_t^2}{\mu} - \frac{\mu}{8} \sum_{t=1}^T (\hat{y}_t - y_t(i^*))^2$$

provided  $\max_i \max_t |\hat{y}_t - y_t(i)| \leq M$ .

*Proof.* First, observe that  $\eta < \frac{1}{2}$  by assumption on  $\mu$ , making it a valid choice. Using Lemma 2 and

$$|c_t(\hat{y}_t - y_t^*)| \leq \frac{c_t^2}{\lambda} + \frac{\lambda}{4}(\hat{y}_t - y_t^*)^2 \quad (16)$$

we obtain

$$\begin{aligned} \sum_{t=1}^T (\hat{y}_t - y_t^*) \ell'_t(\hat{y}_t) &= \sum_{t=1}^T (\hat{y}_t - y_t^*) g_t - \sum_{t=1}^T (\hat{y}_t - y_t^*) c_t \\ &\leq \frac{M^2 \log(K)}{\eta} + \sum_{t=1}^T \left( \eta (\hat{y}_t - y_t(i^*))^2 + \frac{\mu}{4} (\hat{y}_t - y_t^*)^2 + \frac{c_t^2}{\mu} \right). \end{aligned}$$

We conclude by using Lemma 1:

$$\mathcal{R}_T \leq \frac{M^2 \log(K)}{\eta} + \left( \eta + \frac{\mu}{4} - \frac{\mu}{2} \right) \sum_{t=1}^T (\hat{y}_t - y_t(i^*))^2 + \frac{1}{\mu} \sum_{t=1}^T c_t^2,$$

after which replacing  $\eta = \frac{\mu}{8}$  completes the proof.  $\square$

We now restate Theorem 6, after which we prove its result.

**Theorem 6.** Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses and corruptions  $c_1, \dots, c_T \in \mathbb{R}$ . Then the predictions  $\hat{y}_t$  of Algorithm 2 run with inputs  $\eta = \frac{\mu}{8}$ ,  $\sigma = D^2$ ,  $G \geq \max_t |\ell'_t(\hat{y}_t) + c_t|$ ,  $Z > 0$ , feedback  $g_t = \ell'_t(\hat{y}_t) + c_t$ , and  $\kappa_t = 1$ , satisfy

$$\mathcal{R}_T \leq \frac{4dG^2}{\mu} \log \left( 1 + \frac{TD^2\mu^2 \max_t \|\mathbf{x}_t\|_2^2}{2d} \right) + \frac{4G^2}{\mu} + \sum_{t=1}^T \frac{c_t^2}{\mu} - \frac{\mu}{8} \sum_{t=1}^T (\hat{y}_t - y_t^*)^2,$$

for any  $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{R}^d$ , and for any  $\mathbf{u} \in \mathcal{W}_T \equiv \bigcap_{t=1}^T \{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x}_t \rangle| \leq Z\}$  such that  $\|\mathbf{u}\|_2 \leq D$  and  $y_t^* = \langle \mathbf{u}, \mathbf{x}_t \rangle$  for all  $t \geq 1$ .

*Proof.* As in the proof of Theorem 5, using Lemma 3 and equation (16) we obtain

$$\begin{aligned} \sum_{t=1}^T (\hat{y}_t - y_t^*) \ell'_t(\hat{y}_t) &= \sum_{t=1}^T (\hat{y}_t - y_t^*) g_t - \sum_{t=1}^T (\hat{y}_t - y_t^*) c_t \\ &\leq \frac{G^2}{2\eta} + \frac{dG^2}{2\eta} \log \left( 1 + \frac{1}{2} D^2 \mu^2 \left( \max_t \|\mathbf{x}_t\|_2^2 \right) \frac{T}{d} \right) + \sum_{t=1}^T \left( \eta (\hat{y}_t - y_t^*)^2 + \frac{c_t^2}{\mu} \right). \end{aligned}$$

We continue by using Lemma 1:

$$\mathcal{R}_T \leq \frac{G^2}{2\eta} + \frac{dG^2}{4\eta} \log \left( 1 + \frac{1}{2} D^2 \mu^2 \left( \max_t \|\mathbf{x}_t\|_2^2 \right) \frac{T}{d} \right) + \left( \eta + \frac{\mu}{4} - \frac{\mu}{2} \right) \sum_{t=1}^T (\hat{y}_t - y_t^*)^2 + \sum_{t=1}^T \frac{c_t^2}{\mu},$$

after which replacing  $\eta = \frac{\mu}{8}$  completes the proof.  $\square$

## D Details of Section 6 (Selective Sampling)

The following Theorem is a detailed version of Theorem 7

**Theorem 11.** Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses. Then the predictions  $\hat{y}_t$  of Algorithm 1 run with inputs  $M \geq \max_t |\ell'_t(\hat{y}_t)|$ ,  $\eta = \frac{\mu}{4}$ , feedback  $g_t = \frac{o_t}{q_{t-1}} \ell'_t(\hat{y}_t)$ , and  $\kappa_t = q_t$  satisfy

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] \leq \left( \frac{4M^2(\log(K) + 1)}{\mu^{3/2}\beta} \right)^2 + \frac{4M^2(\log(K) + 1)}{\mu},$$

provided  $\max_i \max_t |\hat{y}_t - y_t(i)| \leq M$ .

*Proof.* First observe that  $q_{t-1} \geq q_t$  and thus  $\kappa_t = q_t$  is a valid choice, where we define  $q_0 = q_1$ . By equation (8) we have

$$\begin{aligned} \sum_{t=1}^T \gamma(\hat{y}_t - y_t(i)) \ell'_t(\hat{y}_t) \frac{o_t}{q_{t-1}} &\leq \frac{\log(K)}{q_T} + \gamma^2 \sum_{t=1}^T \kappa_{t-1} (\hat{y}_t - y_t(i))^2 g_t^2 \\ &\leq \frac{\log(K)}{q_T} + \gamma^2 \sum_{t=1}^T \frac{o_t}{q_{t-1}} (\hat{y}_t - y_t(i))^2 \ell'_t(\hat{y}_t)^2 \\ &\leq \frac{\log(K)}{q_T} + \gamma^2 \sum_{t=1}^T \frac{o_t}{q_t} (\hat{y}_t - y_t(i))^2 M^2, \end{aligned}$$

where in the final inequality we used  $q_t \leq q_{t-1}$  and  $|\ell'_t(\hat{y}_t)| \leq M$ . After dividing both sides of the above inequality by  $\gamma = \frac{\eta}{M^2}$  we find

$$\sum_{t=1}^T (\hat{y}_t - y_t(i)) \ell'_t(\hat{y}_t) \frac{o_t}{q_{t-1}} \leq \frac{M^2 \log(K)}{\eta q_T} + \eta \sum_{t=1}^T \frac{o_t}{q_t} (\hat{y}_t - y_t(i))^2. \quad (17)$$

Following the analysis of the clipping trick by Cutkosky [2019], we have that

$$\sum_{t=1}^T (\hat{y}_t - y_t(i)) \ell'_t(\hat{y}_t) \left( \frac{o_t}{q_t} - \frac{o_t}{q_{t-1}} \right) \leq M^2 \sum_{t=1}^T \left( \frac{1}{q_t} - \frac{1}{q_{t-1}} \right) \leq \frac{M^2}{q_T}, \quad (18)$$

where we used that the sum telescopes, that  $q_t \leq q_{t-1}$ , and that  $q_0 = q_1$ . Summing side by side (17) and (18) we find

$$\sum_{t=1}^T (\hat{y}_t - y_t(i)) \ell'_t(\hat{y}_t) \frac{o_t}{q_t} \leq \frac{M^2(\log(K) + 1)}{\eta q_T} + \eta \sum_{t=1}^T \frac{o_t}{q_t} (\hat{y}_t - y_t(i))^2.$$

Since  $\ell_t$  is  $\mu$ -strongly convex we have that

$$\mathbb{E} [\ell_t(\hat{y}_t) - \ell_t(y_t(i))] \leq \mathbb{E} \left[ (\hat{y}_t - y_t(i)) \ell'_t(\hat{y}_t) - \frac{\mu}{2} (\hat{y}_t - y_t(i))^2 \right],$$

and therefore

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t(i))) \right] \leq \mathbb{E} \left[ \frac{M^2(\log(K) + 1)}{\eta q_T} \right] + \left( \eta - \frac{\mu}{2} \right) \sum_{t=1}^T \mathbb{E} [(\hat{y}_t - y_t(i))^2],$$

where we used that  $\mathbb{E}_{t-1}[o_t] = q_t$ . Now, by using  $\min\{a, b\}^{-1} \leq a^{-1} + b^{-1}$  for  $a, b > 0$  and Jensen's inequality we have that

$$\begin{aligned} \mathbb{E} [q_T^{-1}] &= \mathbb{E} \left[ \left( \min \left\{ 1, \beta / \sqrt{\min_i \sum_{t=1}^T (\hat{y}_t - y_t(i))^2} \right\} \right)^{-1} \right] \\ &\leq 1 + \beta^{-1} \sqrt{\mathbb{E} \left[ \min_i \sum_{t=1}^T (\hat{y}_t - y_t(i))^2 \right]}, \end{aligned}$$

and thus

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] &\leq \frac{M^2(\log(K) + 1)}{\eta} \\ &+ \beta^{-1} \sqrt{\mathbb{E} \left[ \min_i \sum_{t=1}^T (\hat{y}_t - y_t(i))^2 \right]} \frac{M^2(\log(K) + 1)}{\eta} - \mathbb{E} \left[ \sum_{t=1}^T \frac{\mu}{4} (\hat{y}_t - y_t^*)^2 \right]. \end{aligned}$$

Using  $ab \leq \frac{a^2\mu}{4} + \frac{b^2}{\mu}$  for  $a, b > 0$  we continue

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] \leq \left( \frac{M^2(\log(K) + 1)}{\eta\beta\sqrt{\mu}} \right)^2 + \frac{M^2(\log(K) + 1)}{\eta}.$$

Using that  $\eta^{-1} = \frac{4}{\mu}$  we arrive at the conclusion of the proof:

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] \leq \left( \frac{4M^2(\log(K) + 1)}{\mu^{3/2}\beta} \right)^2 + \frac{4M^2(\log(K) + 1)}{\mu}.$$

□

## D.1 Selective Sampling for Online Regression

We now extend our selective sampling results to online regression, where we assume that the feature vector  $\mathbf{x}_t$  is revealed to the learner before issuing a prediction. We run Algorithm 2 and, similarly to Section 6, we request the loss at round  $t$  by drawing a Bernoulli variable  $o_t$  of parameter

$$q_t = \min \left\{ 1, \beta \left/ \min_{\mathbf{w} \in \mathcal{W}_t \cap \{\mathbf{w} : \|\mathbf{w}\|_2 \leq D\}} \sum_{s=1}^t (\hat{y}_s - \langle \mathbf{w}, \mathbf{x}_s \rangle)^2 \right. \right\}, \quad (19)$$

for some  $D, \beta > 0$ . This gives the following expected regret guarantee.

**Theorem 12.** *Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses. Then the predictions  $\hat{y}_t$  of Algorithm 2 run with inputs  $\eta = \frac{\mu}{4}$ ,  $\sigma = D^2$ ,  $G \geq \max_t |\ell'_t(\hat{y}_t)|$ ,  $Z = \frac{1}{2}G$ , feedback  $g_t = \frac{o_t}{q_t} \ell'_t(\hat{y}_t)$  and  $\kappa_t = q_t$ , satisfy*

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] \leq \left( \frac{2\zeta_T}{\mu^{3/2}\beta} \right)^2 + \frac{2\zeta_T}{\mu}.$$

for any  $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{R}^d$ , and for any  $\mathbf{u} \in \mathcal{W}_T \equiv \bigcap_{t=1}^T \{\mathbf{w} : |\langle \mathbf{w}, \mathbf{x}_t \rangle| \leq Z\}$  such that  $\|\mathbf{u}\|_2 \leq D$  and  $y_t^* = \langle \mathbf{u}, \mathbf{x}_t \rangle$  for all  $t \geq 1$ , where

$$\zeta_T = G^2 + dG^2 \log \left( 1 + D^2 \max_t \|\mathbf{x}_t\|_2^2 \left( \frac{T}{2d} + \frac{G^2 T^2}{2d\beta^2} \right) \right).$$

*Proof.* Starting from (9) and replacing  $g_t$  by  $\frac{o_t}{q_t} \ell'_t(\hat{y}_t)$  we find

$$\begin{aligned} &\sum_{t=1}^T \frac{o_t}{q_t} \langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t \ell'_t(\hat{y}_t) \rangle \\ &\leq \frac{\|\mathbf{u}\|_2^2}{2\sigma\gamma} + \frac{d}{2\kappa_T\gamma} \log \left( 1 + 2\sigma\gamma^2 \left( \max_t \|\mathbf{x}_t\|_2^2 \right) \frac{T}{d} \right) + \gamma \sum_{t=1}^T \kappa_t \frac{o_t}{q_t^2} (\langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t \ell'_t(\hat{y}_t) \rangle)^2 \\ &\leq \frac{G^2}{2\eta} + \frac{dG^2}{2q_T\eta} \log \left( 1 + D^2 q_t^{-2} \left( \max_t \|\mathbf{x}_t\|_2^2 \right) \frac{T}{2d} \right) + \gamma \sum_{t=1}^T \frac{o_t}{q_t} (\langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t \rangle)^2, \end{aligned}$$



where in the second inequality we used that  $\kappa_t = q_t$ ,  $\|\mathbf{u}\|_2 \leq D$ ,  $\sigma^2 = D^2$ ,  $\gamma = \frac{\eta}{\sigma^2}$ ,  $\eta \leq \frac{1}{2}$  by assumption on  $\mu$ , and  $|\ell'_t(\hat{y}_t)| \leq G$ . Now, using that  $\langle \mathbf{w}_t - \mathbf{u}, \mathbf{x}_t \ell'_t(\hat{y}_t) \rangle = (\hat{y}_t - y_t^*) \ell'_t(\hat{y}_t)$  and by taking the expectation of both sides of the above and using that  $\mathbb{E}_{t-1}[o_t] = q_t$  we find

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^T (\hat{y}_t - y_t^*) \ell'_t(\hat{y}_t) \right] \\ & \leq \mathbb{E} \left[ \frac{G^2}{2q_T \eta} \right] + \mathbb{E} \left[ \frac{dG^2}{2\eta q_T} \log \left( 1 + q_T^{-2} D^2 \max_t \|\mathbf{x}_t\|_2^2 \frac{T}{2d} \right) \right] + \mathbb{E} \left[ \sum_{t=1}^T \eta (\hat{y}_t - y_t^*)^2 \right] \\ & \leq \mathbb{E} \left[ \frac{G^2}{2q_T \eta} \right] + \mathbb{E} \left[ \frac{dG^2}{2\eta q_T} \log \left( 1 + D^2 \max_t \|\mathbf{x}_t\|_2^2 \left( \frac{T}{2d} + \frac{G^2 T^2}{2d\beta^2} \right) \right) \right] + \mathbb{E} \left[ \sum_{t=1}^T \eta (\hat{y}_t - y_t^*)^2 \right] \\ & = \mathbb{E} \left[ \frac{\zeta_T}{2q_T \eta} \right] + \mathbb{E} \left[ \sum_{t=1}^T \eta (\hat{y}_t - y_t^*)^2 \right], \end{aligned}$$

where in the final inequality we used that  $q_T^{-1} \leq 1 + G\sqrt{T}/\beta$  and defined

$$\zeta_T = G^2 + dG^2 \log \left( 1 + D^2 \max_t \|\mathbf{x}_t\|_2^2 \left( \frac{T}{2d} + \frac{G^2 T^2}{2d\beta^2} \right) \right).$$

Since  $\ell_t$  is  $\mu$ -strongly convex we have that

$$\mathbb{E} [\ell_t(\hat{y}_t) - \ell_t(y_t^*)] \leq \mathbb{E} \left[ (\hat{y}_t - y_t^*) \ell'_t(\hat{y}_t) - \frac{\mu}{2} (\hat{y}_t - y_t^*)^2 \right],$$

and therefore

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] & \leq \mathbb{E} \left[ \frac{\zeta_T}{2q_T \eta} \right] + \mathbb{E} \left[ \sum_{t=1}^T \left( \eta - \frac{\mu}{2} \right) (\hat{y}_t - y_t^*)^2 \right] \\ & \leq \mathbb{E} \left[ \frac{\zeta_T}{2q_T \eta} \right] - \mathbb{E} \left[ \sum_{t=1}^T \frac{\mu}{4} (\hat{y}_t - y_t^*)^2 \right], \end{aligned}$$

where we used that  $\eta = \frac{\mu}{4}$ . Now, by using  $\min\{a, b\}^{-1} \leq a^{-1} + b^{-1}$  for  $a, b > 0$  and Jensen's inequality we have that

$$\begin{aligned} \mathbb{E} [q_T^{-1}] & = \mathbb{E} \left[ \left( \min \left\{ 1, \beta / \sqrt{\min_{\mathbf{w} \in \mathcal{W}_T \cap \{\mathbf{w}: \|\mathbf{w}\|_2 \leq D\}} \sum_{t=1}^T (\hat{y}_t - \langle \mathbf{w}, \mathbf{x}_t \rangle)^2} \right\} \right)^{-1} \right] \\ & \leq 1 + \beta^{-1} \sqrt{\mathbb{E} \left[ \min_{\mathbf{w} \in \mathcal{W}_T \cap \{\mathbf{w}: \|\mathbf{w}\|_2 \leq D\}} \sum_{t=1}^T (\hat{y}_t - \langle \mathbf{w}, \mathbf{x}_t \rangle)^2 \right]}, \end{aligned}$$

and thus

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] \\ & \leq \beta^{-1} \sqrt{\mathbb{E} \left[ \min_{\mathbf{w} \in \mathcal{W}_T \cap \{\mathbf{w}: \|\mathbf{w}\|_2 \leq D\}} \sum_{t=1}^T (\hat{y}_t - \langle \mathbf{w}, \mathbf{x}_t \rangle)^2 \right]} \frac{\zeta_T}{2\eta} + \frac{\zeta_T}{2\eta} - \mathbb{E} \left[ \sum_{t=1}^T \frac{\mu}{4} (\hat{y}_t - y_t^*)^2 \right]. \end{aligned}$$

Using  $ab \leq \frac{a^2 \mu}{4} + \frac{b^2}{\mu}$  for  $a, b > 0$  we continue

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] \leq \left( \frac{\zeta_T}{2\eta\beta\sqrt{\mu}} \right)^2 + \frac{\zeta_T}{2\eta}.$$

Using that  $\eta^{-1} = \frac{4}{\mu}$  we arrive at the conclusion of the proof:

$$\mathbb{E} \left[ \sum_{t=1}^T (\ell_t(\hat{y}_t) - \ell_t(y_t^*)) \right] \leq \left( \frac{2\zeta_T}{\mu^{3/2}\beta} \right)^2 + \frac{2\zeta_T}{\mu}.$$

□

---

**Algorithm 4:** AdaHedge with abstention [Van der Hoeven, 2020]

---

**Input** AdaHedge**for**  $t = 1, \dots, T$  **do**    Receive expert predictions  $y_t(1), \dots, y_t(K)$     Obtain distribution  $\mathbf{p}_t$  from AdaHedge    Set  $\hat{y}_t = \sum_{i=1}^K p_t(i) y_t(i)$     Set  $\tilde{y}_t = \text{sign}(\hat{y}_t)$     Set  $b_t = 1 - |\hat{y}_t|$     Set sample  $a_t$  from a Bernoulli distribution with parameter  $1 - b_t$     If  $a_t = 1$ , predict  $\tilde{y}_t$ , otherwise abstain from prediction    Receive  $y_t$ , send  $\frac{1}{2}(1 - y_t(i)y_t)$  as the loss of the  $i$ -th expert to AdaHedge

---

## E Exponentially Weighted Average

Here we provide a brief description of the Exponentially Weighted Average (EWA) algorithm [Vovk, 1990, Littlestone and Warmuth, 1994] on a discrete set of experts. EWA maintains a distribution  $\mathbf{p}_t$  over the experts, where the mass on expert  $i$  is given by

$$p_t(i) \propto \exp\left(-\eta \sum_{s=1}^{t-1} \ell_s(y_s(i))\right),$$

where  $\eta > 0$  is the learning rate. A standard result is that the regret of EWA can be bounded as—see, for example, Van der Hoeven et al. [2018, Lemma 1]:

$$\mathcal{R}_T \leq \frac{\log(K)}{\eta} + \sum_{t=1}^T \left( \ell_t(\hat{y}_t) + \frac{1}{\eta} \log(\mathbb{E}_{i \sim \mathbf{p}_t} [\exp(-\eta \ell_t(y_t(i)))] \right). \quad (20)$$

For  $\alpha$ -exp concave losses, which are losses for which  $g_t(y) = \exp(-\alpha \ell_t(y))$  is concave, we can further bound (20) by choosing  $\eta = \alpha$  and using Jensen's inequality:

$$\mathcal{R}_T \leq \frac{\log(K)}{\alpha}.$$

For  $\mu$ -strongly convex and  $G$ -Lipschitz losses we recover the optimal rate by using the fact that  $\mu$ -strongly convex losses are  $\frac{\mu}{G^2}$ -exp concave [Bubeck, 2011, proposition 1.2]:

$$\mathcal{R}_T \leq \frac{G^2 \log(K)}{\mu}.$$

## F High Probability Regret Bounds for Online Prediction with Abstention

We consider the following generalization of the online learning with abstention setting due to Van der Hoeven [2020]. In each round  $t = 1, \dots, T$  the learner receives expert predictions  $y_t(i) \in [-1, 1]$  and the learner can then either predict  $\tilde{y}_t \in [-1, 1]$  or abstain from prediction. If the learner predicts with  $\tilde{y}_t$  the learner suffers half the hinge loss  $\ell_t(y) = \frac{1}{2}(1 - yy_t)$ , where  $y_t \in \{-1, 1\}$ . If the learner abstains from prediction the learner suffers abstentions cost  $\rho \in [0, \frac{1}{2})$ . Let  $a_t = 1$  if the learner predicts with  $\tilde{y}_t$  and let  $a_t = 0$  if the learner abstains from prediction. The goal is to control the following definition of regret:

$$\sum_{t=1}^T a_t \ell_t(\tilde{y}_t) + (1 - a_t) \rho - \sum_{t=1}^T \ell_t(y_t^*),$$

where  $y_t^* = y_t(i^*)$  and  $i^* = \text{argmin}_i \sum_{t=1}^T \ell_t(y_t(i))$ . In the online prediction with abstention setting an analog of Lemma 1 can be derived. However, in the technical part we instead use Algorithm 2 by Van der Hoeven [2020]. Algorithm 2 by Van der Hoeven [2020] samples  $a_t = 1$  with probability  $1 - b_t$  and  $a_t = 0$  with probability  $b_t$ , where  $b_t = 1 - |\hat{y}_t|$  and  $\hat{y}_t = \sum_{i=1}^K p_t(i) y_t(i)$ . Distributions  $\mathbf{p}_t$  come from AdaHedge [De Rooij et al., 2014] and if  $a_t = 1$  we predict with  $\tilde{y}_t = \text{sign}(\hat{y}_t)$ . Algorithm 2 by Van der Hoeven [2020], or Algorithm 4 in this paper, has the following expected regret guarantee.

**Lemma 7.** For any  $\eta > 0$  Algorithm 4 guarantees

$$\begin{aligned} & \sum_{t=1}^T ((1 - b_t) \ell_t(\tilde{y}_t) + b_t \rho) \\ & \leq \sum_{t=1}^T \ell_t(y_t^*) + \frac{\log(K)}{\eta} + \frac{4}{3} \log(K) + 2 + \sum_{t=1}^T \frac{1}{2} (\eta - (1 - 2\rho)) (1 - |\hat{y}_t|). \end{aligned}$$

*Proof.* From [Van der Hoeven, 2020, Lemma 3] we have that for any  $\eta > 0$

$$\begin{aligned} & \sum_{t=1}^T ((1 - b_t) \ell_t(\text{sign}(\hat{y}_t)) + b_t \rho) \leq \sum_{t=1}^T \ell_t(y_t^*) + \frac{4}{3} \log(K) + 2 \\ & + \frac{\log(K)}{\eta} + \eta \sum_{t=1}^T (\mathbb{E}_{i \sim p_t} [(\ell_t(\hat{y}_t) - \ell_t(y_t(i)))^2] + ((1 - b_t) \ell_t(\text{sign}(\hat{y}_t)) + b_t) - \ell_t(\hat{y}_t)). \end{aligned}$$

Now, by Van der Hoeven [2020, equation (16)] we have that

$$\begin{aligned} & \mathbb{E}_{i \sim p_t} [(\ell_t(\hat{y}_t) - \ell_t(y_t(i)))^2] + ((1 - b_t) \ell_t(\text{sign}(\hat{y}_t)) + b_t) - \ell_t(\hat{y}_t) \\ & \leq \rho(1 - |\hat{y}_t|) + \eta \frac{1}{2} (1 - |\hat{y}_t|) - \frac{1}{2} (1 - |\hat{y}_t|), \end{aligned}$$

and thus

$$\begin{aligned} & \sum_{t=1}^T ((1 - b_t) \ell_t(\text{sign}(\hat{y}_t)) + b_t \rho) \\ & \leq \sum_{t=1}^T \ell_t(y_t^*) + \frac{\log(K)}{\eta} + \frac{4}{3} \log(K) + 2 + \sum_{t=1}^T (\eta \frac{1}{2} - \frac{1}{2} (1 - 2\rho)) (1 - |\hat{y}_t|), \end{aligned}$$

which completes the proof.  $\square$

To see why Lemma 7 is the analog of Lemma 1 for the online prediction with abstention setting observe that by choosing  $\eta < 1 - 2\rho$  we recover a bound akin to (3). In particular, by using Lemma 7, choosing  $\eta = \frac{1}{2}(1 - 2\rho)$  we find

$$\begin{aligned} \sum_{t=1}^T ((1 - b_t) \ell_t(\text{sign}(\hat{y}_t)) + b_t \rho) & \leq \sum_{t=1}^T \ell_t(y_t^*) + \frac{2 \log(K)}{(1 - 2\rho)} + \frac{4}{3} \log(K) + 2 \\ & - \frac{(1 - 2\rho)}{4} \sum_{t=1}^T (1 - |\hat{y}_t|). \end{aligned}$$

meaning we can exploit the negative  $\frac{(1-2\rho)}{4} \sum_{t=1}^T (1 - |\hat{y}_t|)$  in online prediction with abstention in a similar manner as we exploited the negative variance term in online learning with strongly convex losses. As an application of Lemma 7 we provide a high-probability bound for online prediction with abstention.

**Theorem 13.** For  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , Algorithm 4 guarantees

$$\sum_{t=1}^T (a_t \ell_t(\text{sign}(\hat{y}_t)) + (1 - a_t) \rho) \leq \frac{2 \log(K)}{1 - 2\rho} + \frac{21 \log(1/\delta)}{8(1 - 2\rho)} + \frac{4}{3} \log(K) + 2.$$

*Proof.* Let  $r_t = a_t \ell_t(\text{sign}(\hat{y}_t)) + (1 - a_t) \rho - ((1 - b_t) \ell_t(\text{sign}(\hat{y}_t)) + b_t \rho)$ . Since  $\mathbb{E}[a_t \ell_t(\text{sign}(\hat{y}_t)) + (1 - a_t) \rho] = ((1 - b_t) \ell_t(\text{sign}(\hat{y}_t)) + b_t \rho)$  and  $|r_t| \leq 1$ , by Lemma 4, for  $\delta \in (0, 1)$  and  $\lambda \in [0, 1]$ , with probability at least  $1 - \delta$  we have that

$$r_t \leq \frac{\log(1/\delta)}{\lambda} + \lambda \frac{3}{4} \sum_{t=1}^T \mathbb{E}[r_t^2]$$

Now, let us study  $\mathbb{E}[r_t^2] = \mathbb{E}[(a_t \ell_t(\text{sign}(\hat{y}_t)) + (1 - a_t)\rho)^2] - ((1 - b_t)\ell_t(\text{sign}(\hat{y}_t)) + b_t\rho)^2$ . If  $\text{sign}(\hat{y}_t) = y_t$  then

$$\mathbb{E}[r_t^2] = b_t\rho^2 - (b_t\rho)^2 \leq \frac{1}{4}b_t = \frac{1}{4}(1 - |\hat{y}_t|).$$

If  $\text{sign}(\hat{y}_t) \neq y_t$  then

$$\begin{aligned} \mathbb{E}[r_t^2] &= (1 - b_t) + b_t\rho^2 - (1 - b_t + b_t\rho)^2 \\ &= (1 - b_t) + b_t\rho^2 - b_t^2\rho^2 - (1 - b_t)^2 + (1 - b_t)b_t\rho \\ &= |\hat{y}_t| - |\hat{y}_t|^2 + (1 - |\hat{y}_t|)\rho^2 - (1 - |\hat{y}_t|)^2\rho^2 + |\hat{y}_t|(1 - |\hat{y}_t|)\rho \\ &\leq |\hat{y}_t|(1 - |\hat{y}_t|) + (1 - |\hat{y}_t|)\rho^2 \\ &\leq \frac{7}{4}(1 - |\hat{y}_t|), \end{aligned}$$

thus we may conclude that  $\mathbb{E}[r_t^2] \leq \frac{7}{4}(1 - |\hat{y}_t|)$ . Combining the above with Lemma 7 we have that with probability at least  $1 - \delta$

$$\begin{aligned} \sum_{t=1}^T (a_t \ell_t(\text{sign}(\hat{y}_t)) + (1 - a_t)\rho) &\leq \frac{\log(K)}{\eta} + \frac{\log(1/\delta)}{\lambda} + \frac{4}{3}\log(K) + 2 \\ &\quad + \sum_{t=1}^T (\eta \frac{1}{2} + \lambda \frac{21}{16} - \frac{1}{2}(1 - 2\rho)). \end{aligned}$$

Since  $\rho < 1/2$ ,  $\eta = \frac{1}{2}(1 - 2\rho)$  and  $\lambda = \frac{8}{21}(1 - 2\rho)$  are valid choices and we obtain

$$\sum_{t=1}^T (a_t \ell_t(\text{sign}(\hat{y}_t)) + (1 - a_t)\rho) \leq \frac{2\log(K)}{1 - 2\rho} + \frac{21\log(1/\delta)}{8(1 - 2\rho)} + \frac{4}{3}\log(K) + 2,$$

which completes the proof.  $\square$

## G Free Restarts

---

**Algorithm 5:** Restarting

---

**Input**  $\eta > 0, M > 0$

**Initialize**  $\tau_1 = 1, \nu = 1$ , Algorithm 1 with inputs  $M$  and  $\eta$

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

Receive expert predictions  $y_t(1), \dots, y_t(K)$

Send expert predictions to Algorithm 1 and receive  $\hat{y}_t$

Predict  $\hat{y}_t$  and receive loss  $\ell_t$

Send  $g_t = \ell'_t(\hat{y}_t)$  and  $\kappa_t = 1$  to Algorithm 1

**if**  $\min_i \frac{\mu}{4} \sum_{s=\tau_\nu}^t (\hat{y}_s - y_s(i))^2 \geq \frac{4M^2 \log(K)}{\mu}$  **then**

Restart Algorithm 1 with inputs  $M$  and  $\eta$

Set  $\nu = \nu + 1$  and  $\tau_{\nu+1} = t + 1$

Set  $\tau_{\nu+1} = T + 1$

---

In this section, we introduce another way to exploit negative regret. The idea is the following. We keep track of  $\min_i \sum_{s=1}^t (\hat{y}_s - y_s(i))^2$ , which is a lower bound on  $\sum_{s=1}^t (\hat{y}_s - y_s^*)^2$ , and as soon as  $\frac{\mu}{4} \min_i \sum_{s=1}^t (\hat{y}_s - y_s(i))^2 \geq \frac{4M^2 \log(K)}{\mu}$ , Lemma 1 with an appropriate  $\eta$  ensures that  $R_t \leq 0$ . This implies that we may restart the algorithm for free, and compete with a new best expert from that point on. This approach leads to a simplified dynamic regret bound—see, for example, Zhang et al. [2018] or the references therein for a discussion of dynamic regret—in which the expert we are competing against may change in all rounds where the algorithm restarts. Our simplified dynamic regret bound is never larger than the standard regret bound. We denote by  $\nu$  the number of restarts and by  $\tau_\nu$  the first round of restart  $\nu$ . The algorithm can be found in Algorithm 5 and its regret guarantee can be found in Theorem 14 below.

**Theorem 14.** Fix an arbitrary sequence  $\ell_1, \dots, \ell_T$  of  $\mu$ -strongly convex differentiable losses. Then the predictions  $\hat{y}_t$  of Algorithm 5 run with  $g_t = \ell'_t(\hat{y}_t)$ ,  $\kappa_t = 1$ , and inputs  $\eta = \frac{\mu}{4}$  and  $M$  guarantees

$$\sum_{n=1}^{\nu} \max_i \sum_{t=\tau_n}^{\tau_{n+1}-1} \left( \ell_t(\hat{y}_t) - \ell_t(y_t(i)) \right) \leq \frac{4M^2 \log(K)}{\mu},$$

provided that  $\max_t \max_i |y_t(i)| \leq \frac{1}{2}M$  and  $\max_t |\ell'_t(\hat{y}_t)| \leq M$ .

*Proof.* First, observe that  $\eta < \frac{1}{2}$  by assumption on  $\mu$ , making it a valid choice for Algorithm 1. For any  $n < \nu$  we have that for any  $i \in [K]$ , by Lemma 2 and Lemma 1

$$\sum_{t=\tau_n}^{\tau_{n+1}-1} \left( \ell_t(\hat{y}_t) - \ell_t(y_t(i)) \right) \leq \frac{4M^2 \log(K)}{\mu} - \frac{\mu}{4} \sum_{t=\tau_n}^{\tau_{n+1}-1} (\hat{y}_t - y_t(i))^2.$$

Since  $n < \nu$  we must have that

$$\frac{4M^2 \log(K)}{\mu} \leq \min_j \frac{\mu}{4} \sum_{t=\tau_n}^{\tau_{n+1}-1} (\hat{y}_t - y_t(j))^2 \leq \frac{\mu}{4} \sum_{t=\tau_n}^{\tau_{n+1}-1} (\hat{y}_t - y_t(i))^2$$

and thus

$$\sum_{t=\tau_n}^{\tau_{n+1}-1} \left( \ell_t(\hat{y}_t) - \ell_t(y_t(i)) \right) \leq 0.$$

For  $n = \nu$  we have that for any  $i \in [K]$ , by Lemma 2 and Lemma 1

$$\sum_{t=\tau_n}^{\tau_{n+1}-1} \left( \ell_t(\hat{y}_t) - \ell_t(y_t(i)) \right) \leq \frac{4M^2 \log(K)}{\mu}$$

which completes the proof. □