

# Fast Rates for Online Prediction with Abstention

**Gergely Neu**

*Universitat Pompeu Fabra, Barcelona, Spain*

**Nikita Zhivotovskiy**

*Google Research, Brain Team, Zürich, Switzerland*

GERGELY.NEU@GMAIL.COM

ZHIVOTOVSKIY@GOOGLE.COM

**Editors:** Jacob Abernethy and Shivani Agarwal

## Abstract

In the setting of sequential prediction of individual  $\{0, 1\}$ -sequences with expert advice, we show that by allowing the learner to abstain from the prediction by paying a cost marginally smaller than  $\frac{1}{2}$  (say, 0.49), it is possible to achieve expected regret bounds that are independent of the time horizon  $T$ . We exactly characterize the dependence on the abstention cost  $c$  and the number of experts  $N$  by providing matching upper and lower bounds of order  $\frac{\log N}{1-2c}$ , which is to be contrasted with the best possible rate of  $\sqrt{T \log N}$  that is available without the option to abstain. We also discuss various extensions of our model, including a setting where the sequence of abstention costs can change arbitrarily over time, where we show regret bounds interpolating between the slow and the fast rates mentioned above, under some natural assumptions on the sequence of abstention costs.

## 1. Introduction

Consider the problem of online prediction of individual sequences which is one of the first and well-studied *online learning* models (Cesa-Bianchi and Lugosi, 2006). In this setup, a *learner* and an *environment* interact in a sequence of rounds  $t = 1, \dots, T$  as follows. In round  $t$ , the learner observes the  $\{0, 1\}$ -valued predictions of  $N$  *experts*, denoted as  $y_{t,1}, y_{t,2}, \dots, y_{t,N}$ . Based on these observations and possibly some use of randomness, the learner predicts  $\hat{y}_t \in \{0, 1\}$  and then the environment reveals  $y_t \in \{0, 1\}$ . In particular, the environment can be aware of the strategy of the learner but not the random bits used for producing  $\hat{y}_t$ . Having made its prediction, the learner suffers the real-valued loss  $\hat{\ell}_t = \ell(\hat{y}_t, y_t)$ , which in this paper is chosen as the binary loss  $\hat{\ell}_t = \mathbb{I}_{\{\hat{y}_t \neq y_t\}}$ . The aim of the learner is to minimize its *regret* which is the difference between the total loss of the prediction strategy and the total loss of the best fixed expert chosen in full knowledge of the sequence of outcomes.

Since the naive upper bound  $O(T)$  always holds for the regret, we are aiming for better dependencies. It is well known that it is not possible to guarantee non-trivial bounds on the regret when the learner’s predictions are deterministic. Therefore, we allow the learner to randomize its decisions and measure its performance in terms of the *expected regret* defined formally as

$$R_T = \max_{i \in [N]} \mathbb{E} \left[ \sum_{t=1}^T (\hat{\ell}_t - \ell_{t,i}) \right],$$

where the expectation is taken with respect to the randomness injected by the learner and  $\ell_{t,i}$  is the loss suffered by expert  $i$  in round  $t$ . As shown in the seminal work of Littlestone and Warmuth (1994), it is possible to construct an algorithm that achieves regret of order  $O(\sqrt{T \log N})$  in this

setting, and this is essentially the best one we hope for as long as the binary loss together with randomization is considered—see also the classic works of Hannan (1957), Cover (1965) and Vovk (1990) and the excellent monograph by Cesa-Bianchi and Lugosi (2006). Notably, the result holds without any assumptions on the behaviour of the environment which can be completely adversarial.

One natural direction of research is to understand when it is possible to improve on the worst-case regret guarantees mentioned above. Of particular interest are bounds that are independent on the time horizon  $T$ . Bounds of this type are sometimes referred to as the *fast rates* in online learning (to be contrasted with the worst-case *slow rate* mentioned above). There are several conditions on the loss function or the sequence of outcomes that are known to imply fast rates. A classic result due to Haussler, Kivinen, and Warmuth (1998) and Vovk (1990) claims that if the learner is allowed to output  $y_t \in [0, 1]$  and the loss function  $\ell$  is convex with sufficient curvature, then there exists a deterministic prediction strategy guaranteeing  $R_T = O(\log N)$ . A second type of assumptions that can lead to fast rates even in the case of the binary loss is that the outcomes  $y_t$  are generated in an i.i.d. manner, rather than being chosen by an adversarial environment. In this favourable setup, assuming that the set of experts and the distribution of  $y_t$  satisfy the so-called *Bernstein assumption* (introduced originally in the statistical learning literature by Bartlett and Mendelson, 2006), Koolen et al. (2016) show that it is possible to obtain intermediate rates between  $O(\log N)$  and  $O(\sqrt{T \log N})$  depending on the parameters in the Bernstein assumption. Similar results are well known in the statistical learning setup (see, e.g., Tsybakov, 2004). For an extensive survey on fast rates in online learning, we refer the interested reader to van Erven et al. (2015).

In this paper we show that a simple variation on the basic protocol of online binary prediction also permits regret of  $O(\log N)$ : instead of forcing the learner to output a prediction in  $\{0, 1\}$ , we also allow the learner to *abstain from prediction*. This model is sometimes referred to as *Chow’s reject option model*, introduced in the seminal work of Chow (1970). This setup was intensively studied in *statistical learning* where the learner can output one of three values  $\{0, 1, *\}$  and the value  $*$  corresponds to an abstention (see, e.g., the works of Herbei and Wegkamp, 2006; Bartlett and Wegkamp, 2008; Bousquet and Zhivotovskiy, 2019). In this model the loss is assumed to be binary on  $\{0, 1\}$  outputs but the price for abstention is assumed to be equal to  $c < \frac{1}{2}$ . This means that the learner may profit from using the option to abstain whenever they believe that there is no way to predict  $y_t$  better than by flipping an unbiased coin. We only assume that  $c$  is marginally smaller than  $\frac{1}{2}$  (say, 0.49), so that there is no reason for the learner to abstain too often if they want to have a small regret.

#### The protocol of online binary prediction with abstentions

For each round  $t = 1, 2, \dots, T$ , repeat:

1. The learner observes the  $\{0, 1\}$ -valued predictions of  $N$  experts, denoted as  $y_{t,1}, \dots, y_{t,N}$ ,
2. based on these observations and possibly some use of randomness, the learner predicts  $\hat{y}_t \in \{0, *, 1\}$ ,
3. the environment reveals  $y_t \in \{0, 1\}$ ,
4. the learner suffers the loss  $\hat{\ell}_t = \mathbb{I}_{\{\hat{y}_t \neq y_t\}}$  if  $\hat{y}_t \in \{0, 1\}$ , otherwise a loss of  $\hat{\ell}_t = c \leq \frac{1}{2}$  if  $\hat{y}_t = *$ .

We stress that even when deciding to abstain, the learner still gets to observe the true outcome. The first thing to notice regarding the hardness of our setting is that the learner still needs to randomize its decisions in order to achieve sublinear regret. Specifically, one can easily adapt the classic counterexample of Cover (1965) to show that any deterministic strategy will result in a regret of at least  $\frac{cT}{2}$  for the learner—we provide the details in Appendix A.

We are not the first to explore the effects of allowing abstentions in the context of online learning. The most notable contribution is the work of Cortes, DeSalvo, Gentile, Mohri, and Yang (2018), who extended the model of Herbei and Wegkamp (2006); Bartlett and Wegkamp (2008) from the setting of statistical learning to that of online learning and provided a range of results for both stochastic and adversarial environments. A crucial feature of their setup is that abstention is an action available to the *experts*: each expert produces a prediction in  $\{0, *, 1\}$ , which is then used by the learner for producing its their own randomized action. Furthermore, when deciding to abstain, the learner does not get to observe the outcome produced by the environment. In contrast, our setting treats abstention as an action *only available to the learner*, and the true outcome is always revealed to the learner at the end of the round independently of the learner’s decision. Thus, the setting of Cortes et al. (2018) is in many senses more complicated than ours, and as a result their regret guarantees are weaker: they only match the worst-case rate of  $O(\sqrt{T \log N})$ . It is unclear if it is possible to refine their techniques to obtain fast rates comparable to the ones we prove. Another work related to ours is that of Zhang and Chaudhuri (2016), who considered an online prediction problem with an abstention option under a *realizability assumption* and noted the possibility of adapting Chow’s reject option model to the sequential prediction setting as an interesting direction of future work—which is precisely what we address in this paper.

Our pursuit of fast rates is motivated by the recent results of Bousquet and Zhivotovskiy (2019) who have shown for the first time that allowing the learner to abstain from prediction enables fast rates in *statistical learning*. Notably, their results were obtained without any further assumptions on the data distribution (i.e., no Bernstein condition) or on the curvature of the loss function (i.e., binary loss). Their approach is based on an application of the empirical risk minimization principle on an extended set of classifiers, resulting in a *deterministic* algorithm guaranteeing an excess risk of at most  $O(\frac{\log N}{T})$  when the set of classifiers is finite. While our results are of similar flavor, the approach that we put forward in this paper is different in almost every respect from theirs.

**Our contributions.** Our main contribution the following: For a fixed abstention cost  $c < \frac{1}{2}$  we show in Section 2 that there is an algorithm with the regret bound

$$R_T \leq \frac{\log N}{1 - 2c} \wedge \sqrt{\frac{T \log N}{2}}.$$

We also prove that our bound is tight up to multiplicative constant factors if  $c$  is bounded away from zero (Section 4). Assuming that the abstention cost  $c_t$  changes in time, we prove intermediate rates between  $O(\log N)$  and  $O(\sqrt{T \log N})$  in Section 3. Finally, we provide two natural extensions of our main result in Section 5.

**Notation.** We define  $a \wedge b = \min\{a, b\}$  and  $a \vee b = \max\{a, b\}$ . We also use the standard  $O(\cdot), \Theta(\cdot), \Omega(\cdot)$  notation. The indicator of the event  $A$  is denoted by  $\mathbb{I}_{\{A\}}$ . Given integer  $K$  we define  $[K] = \{1, \dots, K\}$ .

## 2. Algorithm and main result

Our algorithm is based on the classic exponentially weighted forecaster of [Littlestone and Warmuth \(1994\)](#) (see also [Vovk, 1990](#); [Freund and Schapire, 1997](#)), and is described as follows. In each round  $t$ , the algorithm computes the set of weights

$$w_{t,i} = e^{-\eta \sum_{k=1}^{t-1} \ell_{k,i}}$$

for some learning rate  $\eta > 0$ , the corresponding probability distribution over all experts  $i \in [N]$

$$q_{t,i} = \frac{w_{t,i}}{\sum_j w_{t,j}}, \quad \text{and the aggregated prediction} \quad p_t = \sum_i q_{t,i} y_{t,i}.$$

The crucial step in the algorithm design is generating a randomized  $\{0, *, 1\}$ -valued prediction based on  $p_t$ . The key intuition is that the learner should abstain whenever  $p_t = \frac{1}{2}$  and output a deterministic prediction whenever  $p_t \in \{0, 1\}$ . We choose to linearly interpolate between these extreme cases by producing a randomized output supported on  $\{0, *\}$  when  $p_t < \frac{1}{2}$  and  $\{*, 1\}$  otherwise. Precisely, our algorithm will base its decision on the probability of the most likely label  $p_t^* = p_t \vee (1 - p_t)$ : it will abstain with probability  $\alpha_t = 2(1 - p_t^*)$ , and output the most likely label with probability  $1 - \alpha_t$ . The procedure is compactly presented as [Algorithm 1](#).

---

**Algorithm 1** Online prediction with abstentions.

---

**Input:** learning rate  $\eta > 0$ .

**Initialization:** set  $w_{1,i} = 1$  for all  $i \in [N]$ .

**For**  $t = 1, 2, \dots, T$ , **repeat**

1. Observe the advice  $y_{t,i}$  of all experts  $i \in [N]$ .
  2. Calculate the mean prediction  $p_t = \sum_{i \in [N]} \frac{w_{t,i} y_{t,i}}{\sum_{j \in [N]} w_{t,j}}$ .
  3. Let  $p_t^* = p_t \vee (1 - p_t)$  and  $k_t^* = \mathbb{I}_{\{p_t \geq \frac{1}{2}\}}$ , that is, the label attaining the highest probability.
  4. Let  $\alpha_t = 2(1 - p_t^*)$  and predict  $*$  with probability  $\alpha_t$  and  $k_t^*$  with probability  $1 - \alpha_t$ .
  5. Observe the label  $y_t$  and update the weights as  $w_{t+1,i} = w_{t,i} e^{-\eta \ell_{t,i}}$ .
- 

Note that the abstention probability  $\alpha_t$  is in  $[0, 1]$  due to the fact that  $p_t^* \geq \frac{1}{2}$ . To gain some intuition about this rule, observe that when assigning the numerical value  $* = \frac{1}{2}$  to abstention, the resulting prediction has expectation  $p_t$ . Our main result regarding the performance of our algorithm is the following:

**Theorem 1** *Suppose that  $c < \frac{1}{2}$  and  $\eta \leq 2(1 - 2c)$ . Then, the expected regret of [Algorithm 1](#) satisfies*

$$R_T \leq \frac{\log N}{\eta}.$$

Notably, this theorem shows that whenever the abstention cost  $c$  is bounded away from  $1/2$ , we can set  $\eta = 2(1 - 2c)$  and our algorithm achieves a regret bound that is independent of the time horizon

$T$ . However, when  $c$  is very close to  $1/2$ , one may favor to fall back to the standard worst-case regret guarantee of order  $\sqrt{T \log N}$ . This is, however, easily achieved by choosing a conservative value of  $\eta$  for this unfavorable case. The following simple corollary of Theorem 1 summarizes the rates achieved by our algorithm in all regimes.

**Corollary 2** *Setting  $\eta = 2(1 - 2c) \vee \sqrt{\frac{8 \log N}{T}}$ , the regret of our algorithm satisfies*

$$R_T \leq \frac{\log N}{2(1 - 2c)} \wedge \sqrt{\frac{T \log N}{2}}.$$

We provide the simple proofs of both Theorem 1 and Corollary 2 below. A central object in our analysis is the quantity

$$\tilde{\ell}_t = -\frac{1}{\eta} \log \sum_i q_{t,i} e^{-\eta \ell_{t,i}}, \quad (1)$$

often called the *mix loss* in the literature (Vovk, 1990, 1998; De Rooij et al., 2014). The following classic result highlights the key role of the mix loss in the analysis of the exponentially weighted forecaster:

**Lemma 3** *For any  $\eta > 0$  the cumulative mix loss of the exponentially weighted forecaster satisfies*

$$\sum_{t=1}^T \tilde{\ell}_t \leq \min_i \sum_{t=1}^T \ell_{t,i} + \frac{\log N}{\eta}.$$

We include the proof in Appendix B for the sake of completeness. The heart of our analysis is the following lemma that establishes a connection between the mix loss and the loss suffered by our algorithm:

**Lemma 4** *Suppose that  $c < \frac{1}{2}$  and  $\eta \leq 2(1 - 2c)$ . Then,  $\mathbb{E}[\widehat{\ell}_t] \leq \tilde{\ell}_t$ .*

**Proof** Let us define the misclassification probability of the plain exponentially weighted forecaster as  $r_t = p_t \mathbb{I}_{\{y_t=0\}} + (1 - p_t) \mathbb{I}_{\{y_t=1\}} = \sum_{i=1}^N q_{t,i} \mathbb{I}_{\{y_{t,i} \neq y_t\}}$ . With this notation, the mix loss can be written as

$$-\frac{1}{\eta} \log \sum_i q_{t,i} e^{-\eta \ell_{t,i}} = -\frac{1}{\eta} \log \sum_i q_{t,i} \left( 1 + (e^{-\eta} - 1) \mathbb{I}_{\{y_{t,i} \neq y_t\}} \right) = -\frac{1}{\eta} \log (1 + r_t (e^{-\eta} - 1)).$$

On the other hand, the expected loss of our algorithm is

$$\mathbb{E}[\widehat{\ell}_t] = \alpha_t c + (1 - \alpha_t) \mathbb{I}_{\{k_t^* \neq y_t\}} = 2(1 - p_t^*)c + (2p_t^* - 1) \mathbb{I}_{\{k_t^* \neq y_t\}}.$$

In the case  $y_t = k_t^*$ , we have  $r_t = 1 - p_t^* \leq \frac{1}{2}$  so the above becomes

$$\mathbb{E}[\widehat{\ell}_t] = 2(1 - p_t^*)c = 2r_t c = r_t - (1 - 2c)r_t.$$

Otherwise, we have  $r_t = p_t^* \geq \frac{1}{2}$  so that the expected loss is

$$\mathbb{E}[\widehat{\ell}_t] = 2(1 - p_t^*)c + 2p_t^* - 1 = 2(1 - r_t)c + 2r_t - 1 = r_t - (1 - 2c)(1 - r_t).$$

Also noticing that  $r_t \leq \frac{1}{2}$  holds if and only if  $y_t = k_t^*$ , we conclude that

$$\mathbb{E}[\widehat{\ell}_t] = r_t - (1 - 2c)(r_t \wedge (1 - r_t)).$$

To finish the proof, we define the functions

$$f(r) = -\frac{1}{\eta} \log(1 + r(e^{-\eta} - 1)) \quad \text{and} \quad g(r) = r - (1 - 2c)(r \wedge (1 - r)), \quad (2)$$

and note that  $\widetilde{\ell}_t = f(r_t)$  and  $\mathbb{E}[\widehat{\ell}_t] = g(r_t)$ . We will show that  $g(r) \leq f(r)$  for all  $r \in [0, 1]$ , which will imply the statement of the theorem. To this end, note that both functions are convex and equal to each other at  $r = 0$  and  $r = 1$ , so the desired inequality will hold if the respective derivatives at these two points satisfy  $g'(0) \leq f'(0)$  and  $f'(1) \leq g'(1)$ . To verify this, first observe that

$$f'(r) = \frac{1 - e^{-\eta}}{\eta(1 + r(e^{-\eta} - 1))} = \begin{cases} \frac{1 - e^{-\eta}}{\eta} & \text{if } r = 0, \\ \frac{e^{\eta} - 1}{\eta} & \text{if } r = 1. \end{cases}$$

On the other hand, we have

$$g'(r) = \begin{cases} 2c & \text{if } r = 0, \\ 2(1 - c) & \text{if } r = 1. \end{cases}$$

Thus, the condition on the derivatives we seek is satisfied when

$$\frac{1 - e^{-\eta}}{\eta} \geq 2c \quad \text{and} \quad \frac{e^{\eta} - 1}{\eta} \leq 2(1 - c).$$

Also observing that  $\frac{1 - e^{-\eta}}{\eta} \geq 1 - \frac{\eta}{2}$  and  $\frac{e^{\eta} - 1}{\eta} \leq 1 + \frac{\eta}{2}$  both hold for  $\eta > 0$ , we can verify that the conditions above are satisfied whenever  $\eta \leq 2(1 - 2c)$ , as required in the statement of the theorem. This concludes the proof.  $\blacksquare$

Putting Lemmas 3 and 4 together proves Theorem 1. Corollary 2 is proved by first observing that

$$\mathbb{E}[\widehat{\ell}_t] = r_t - (1 - 2c)(r_t \wedge (1 - r_t)) \leq r_t = \sum_i q_{t,i} \ell_{t,i},$$

holds as long as  $c \leq \frac{1}{2}$ , and then using Hoeffding's lemma (Lemma 2.2 in [Cesa-Bianchi and Lugosi, 2006](#)) that guarantees

$$\sum_i q_{t,i} \ell_{t,i} \leq \widetilde{\ell}_t + \frac{\eta}{8}.$$

Putting this inequality together with Lemma 3, we retrieve the standard bound of the exponentially weighted forecaster:  $R_T \leq \frac{\log N}{\eta} + \frac{\eta T}{8}$ . Combining this bound with the one of Theorem 1 and tuning  $\eta$  proves the corollary.

### 3. Changing abstention costs

Let us now consider a slight variation on our problem where the abstention cost  $c$  may depend on time: in round  $t$ , the cost of abstention is  $c_t \leq \frac{1}{2}$ . The main question we address in this section is whether it is possible to attain fast rates even in the case where  $c_t$  can get arbitrarily close to  $\frac{1}{2}$  in a small number of rounds. It is easy to see that our algorithm given in Section 2 satisfies the following regret bound with its proof presented in Appendix B.

**Proposition 5** *Suppose that  $c_t \leq \frac{1}{2}$  for all  $t$  and  $\eta > 0$ . Then, the regret of our algorithm satisfies*

$$R_T \leq \frac{\log N}{\eta} + \frac{\eta}{8} \sum_{t=1}^T \mathbb{I}_{\{2(1-2c_t) < \eta\}}.$$

Notably, this result asserts that it is not necessary to assume that  $c_t$  is strictly bounded away from  $\frac{1}{2}$  in order to obtain an improvement over the worst-case regret bound. In fact, we show that a clear improvement is possible when making the following quantitative assumption about the behavior of the abstention costs around  $\frac{1}{2}$ :

**Definition 6 (Tsybakov’s condition for abstention costs)** *Assume that there are constants  $\beta > 0, \alpha \in [0, 1)$  such that for any  $x > 0$ ,*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{1/2 - c_t < x\}} \leq \beta x^{\frac{\alpha}{1-\alpha}}. \quad (3)$$

As the name suggests, this condition is directly inspired by the renowned Tsybakov’s margin assumption introduced first by [Mammen and Tsybakov \(1999\)](#) in the context of binary classification and which was further analyzed by [Tsybakov \(2004\)](#). In the context of classification where one has to predict a random label  $Y \in \{0, 1\}$  based on a random instance  $X$ , Tsybakov’s assumption is used to characterize the distribution of the regression function  $\mathbb{E}[Y|X]$  around  $\frac{1}{2}$ , which is related to the hardness of the classification problem at hand. Indeed, instances such that the regression function is further away from  $\frac{1}{2}$  are “easier” to classify than ones closer to this value, so restricting the density of  $\mathbb{E}[Y|X]$  around this problematic region can make the overall learning problem easier—for more details, and relations of Tsybakov’s assumption with the so-called Bernstein assumption we refer the interested reader to [van Erven et al. \(2015\)](#). Similarly, our own online prediction problem becomes “harder” when the abstention costs are densely distributed around  $\frac{1}{2}$ . By analogy with the case of binary classification, we choose to quantify the density of the quantity of interest around  $\frac{1}{2}$  through Tsybakov’s assumption: larger values of  $\alpha$  correspond to less frequent values around the margin, whereas smaller values correspond to more problematic points. Similarly, the situation where the abstention cost is strictly bounded away from  $1/2$  can be seen as an analog of Massart’s margin assumption in the statistical learning setup ([Massart and Nédélec, 2006](#)).

Assuming that the costs satisfy the inequality (3), our algorithm can be easily seen to achieve the following result:

**Corollary 7** *Suppose that the sequence of abstention costs satisfies Tsybakov’s margin condition. Then, setting  $\eta = \left(\frac{\log N}{T}\right)^{\frac{1-\alpha}{2-\alpha}}$ , the regret of our algorithm satisfies*

$$R_T = O\left((\log N)^{\frac{1}{2-\alpha}} T^{\frac{1-\alpha}{2-\alpha}}\right). \quad (4)$$

The proof is immediate given Proposition 5 and the definition of Tsybakov’s condition. We remark that the bound of the very same form as (4) has been recently shown in [Koolen et al. \(2016\)](#) under the assumption that the losses are i.i.d. and that the Bernstein assumption is satisfied. However, our assumption is different and does not imply the Bernstein assumption, even if it leads to similar regret bounds.

One downside of Corollary 7 is that it relies on a choice of the learning rate  $\eta$  that requires prior knowledge of the sequence of abstention costs. More generally, the learning rate  $\eta$  that minimizes the bound of Proposition 5 is clearly a function of the abstention costs. One may wonder if it is possible to attain a regret guarantee comparable to

$$R_T^* = \min_{\eta > 0} \left( \frac{\log N}{\eta} + \frac{\eta}{8} \sum_{t=1}^T \mathbb{I}_{\{2(1-2c_t) < \eta\}} \right),$$

without having prior knowledge of the sequence of abstention costs.

We answer this question in the positive by considering a simple variation of our algorithm based on the exponentially weighted forecaster with adaptive learning rates. Precisely, in each round  $t = 1, 2, \dots, T$ , we will choose a positive learning-rate parameter  $\eta_t$ , compute the weights  $w_{t,i} = e^{-\eta_t \sum_{k=1}^{t-1} \ell_{k,i}}$ , and then use these weights in the same way as our basic algorithm did. The intuition driving our algorithm design is to keep the learning rate as large as possible, and only decrease it when observing high abstention costs. Specifically, we let  $d_t = \sum_{i=1}^t \mathbb{I}_{\{\eta_i \geq 2(1-2c_i)\}}$  be the number of times that the learning rate has exceeded  $2(1-2c_i)$  before round  $t$ , and define our learning rate as  $\eta_{t+1} = \sqrt{\frac{\log N}{d_t}} \wedge 1$ , with  $d_0$  defined as 1. The full algorithm is shown below.

---

**Algorithm 2** Adaptive online prediction with abstentions.

---

**Initialization:** set  $w_{1,i} = 1$  for all  $i \in [N]$ ,  $d_0 = 1$ , and  $\eta_1 = 1$ .

**For**  $t = 1, 2, \dots, T$ , **repeat**

1. Observe the advice  $y_{t,i}$  of all experts  $i \in [N]$ .
  2. Calculate the mean prediction  $p_t = \sum_{i \in [N]} \frac{w_{t,i} y_{t,i}}{\sum_{j \in [N]} w_{t,j}}$ .
  3. Let  $p_t^* = p_t \vee (1 - p_t)$  and  $k_t^* = \mathbb{I}_{\{p_t \geq \frac{1}{2}\}}$ , that is, the label attaining the highest probability.
  4. Let  $\alpha_t = 2(1 - p_t^*)$  and predict  $*$  with probability  $\alpha_t$  and  $k_t^*$  with probability  $1 - \alpha_t$ .
  5. Observe the abstention cost  $c_t$  and if  $\eta_t \geq 1 - 2c_t$ , update  $d_t = d_{t-1} + 1$ .
  6. Set  $\eta_{t+1} = \sqrt{\frac{\log N}{d_t}} \wedge 1$  and compute the weights  $w_{t+1,i} = e^{-\eta_{t+1} \sum_{k=1}^t \ell_{k,i}}$ .
- 

The following theorem (proved in Appendix B) establishes a regret bound for this algorithm.

**Theorem 8** *The regret of Algorithm 2 satisfies  $R_T \leq \frac{15}{8} R_T^* + \frac{5}{4} \sqrt{\log N}$ .*

#### 4. Lower bound

In this section we prove that the regret  $\Theta\left(\frac{\log N}{1-2c} \wedge \sqrt{T \log N}\right)$  is optimal in the fixed abstention cost setup. Precisely, the main result we present here is the following non-asymptotic lower bound which shows that the bound of Theorem 1 is sharp with respect to  $N, c$  and  $T$  provided that  $c$  is separated away from zero:



**Theorem 9** Fix the abstention price  $c \in [\frac{1}{4}, \frac{1}{2}]$ ,  $N \geq 2$  and  $T \geq 4 \log N$ . There is a set of experts of size  $N$  such that for any randomized online prediction strategy  $\hat{y}_t \in \{0, *, 1\}$  there is a strategy of the environment such that

$$R_T = \Omega \left( \frac{\log N}{1 - 2c} \wedge \sqrt{T \log N} \right).$$

The proof of this result is based on observing that any randomized online algorithm can be converted to a randomized batch algorithm in a way that the excess risk in the latter setting is closely related to the regret in the former setting. Thus, any lower bound in the batch setting will imply a lower bound on the best achievable regret. An additional crucially important ingredient is relating the cost of abstention to a suitably chosen strongly convex loss function.

First, we recall the standard statistical learning setup. In this setting we consider the instance space  $\mathcal{X}$  and the label space  $\mathcal{Y} = [-B, B]$ , where  $B > 0$  is a constant. Let  $((X_1, Y_1), \dots, (X_T, Y_T))$  be an i.i.d. sequence of points sampled according to some unknown distribution  $P$  on  $\mathcal{X} \times \mathcal{Y}$ . We will refer to this sequence as *the learning sample* and to this setup as *the batch setting*. Given the loss function  $\ell : \mathcal{Y}^2 \rightarrow \mathbb{R}_+$ , the *risk* of a hypothesis  $g : \mathcal{X} \rightarrow \mathcal{Y}$  is  $\mathbb{E} [\ell(g(X), Y)]$ , where the expectation is taken with respect to  $P$ . We denote  $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$  and  $(Z_1, \dots, Z_T) = ((X_1, Y_1), \dots, (X_T, Y_T))$ . We need the following standard result which can be found in, e.g., [Audibert \(2009\)](#).

**Lemma 10 (Online to batch conversion)** Consider any randomized online learning algorithm  $\mathcal{A}$  that produces the possibly randomized <sup>1</sup> hypothesis  $h_t : \mathcal{X} \rightarrow \mathcal{Y}$  at time  $t$  when the environment uses an i.i.d. sequence  $Z_1, \dots, Z_T$ . We define the randomized batch algorithm  $\mathcal{L}$  that, given the data  $Z_1, \dots, Z_T$  and any new observation  $X$ , predicts its label according to the hypothesis chosen uniformly at random among  $\{h_1, h_2, \dots, h_T\}$ . We denote this randomized hypothesis by  $\hat{g}_T$ . Then, the risk of  $\mathcal{L}$  satisfies

$$\mathbb{E} [\ell(\hat{g}_T(X), Y)] = \frac{1}{T} \mathbb{E} \left[ \sum_{t=1}^T \ell(h_t(X_t), Y_t) \right],$$

where the expectation on the right-hand side integrates over the random hypothesis chosen by  $\mathcal{L}$ , the instances  $Z_1, \dots, Z_T$ , and the randomization of the online algorithm.

Our main technical tool is the following lower bound in the batch setting.

**Lemma 11 (Theorem 8.4 in [Audibert \(2009\)](#))** Let  $B > 0$ ,  $N \geq 2$  and  $q > 1 + \sqrt{\frac{\lfloor \log_2 N \rfloor}{4T}} \wedge 1$ , if  $\mathcal{X}$  contains at least  $\lfloor \log_2(2N) \rfloor$  points, there exists a set  $G$  of  $N$  hypotheses satisfying for any (possibly randomized <sup>2</sup>) estimator  $\hat{g}$  there exists a probability distribution  $P$  on  $\mathcal{Z}$  such that

$$\mathbb{E} [|Y - \hat{g}(X)|^q] - \min_{g \in G} \mathbb{E} [|Y - g(X)|^q] \geq \left( \frac{q}{90(q-1)} \vee e^{-1} \right) B^q \left( \frac{\lfloor \log_2 N \rfloor}{T+1} \wedge 1 \right),$$

where the expectation in  $\mathbb{E} [|Y - \hat{g}(X)|^q]$  is taken with respect to the learning sample,  $(X, Y)$  and the (possible) randomness of the algorithm.

- 
1. To simplify the notation we do not write the parameters responsible for randomization explicitly in the functions.
  2. We note that the original version of this result does not highlight that  $\hat{g}$  is potentially randomized. Nevertheless, the standard argument (see, e.g., the discussions on Yao's minimax principle in [Shamir \(2015\)](#) or Remark 8.1 in [Audibert \(2009\)](#) claiming that any lower bound for the batch setting leads to a lower bound in the sequential prediction setting) shows that Lemma 11 continues to hold in the same form.

With these results in mind we are ready to provide some intuition on the proof of Theorem 9. For  $B = \frac{1}{2}$  it appears that the construction in Lemma 11 uses only the distributions  $P$  such that  $Y = \pm\frac{1}{2}$  and the set  $G$  of hypotheses taking their values in  $\{\pm\frac{1}{2}\}$ . Using Lemma 10 based on our online learning algorithm we may construct an estimator  $\hat{g}$  taking its values in  $\{\pm\frac{1}{2}, 0\}$  where the value 0 corresponds to an abstention. Indeed, it is clear that as long as we consider only  $\pm\frac{1}{2}$  outputs the  $\ell_q$  loss is equivalent to binary. However, if  $\hat{g}$  predicts 0 then for  $q = \log_2 \frac{1}{c}$  the loss will be equal to  $c$  for any  $Y$  as above. The full proof of Theorem 9 is presented in Appendix B.

## 5. Extensions

In this section we provide two natural extensions of our results to the setting of multiclass classification, and the case of binary classification with an infinite set of experts.

### 5.1. Multiclass classification

We now turn our attention to a generalization of our online prediction setup that allows multiple classes. Specifically, we will now consider the case where the outcomes can take  $K$  different values, and the learner's prediction is evaluated through the binary loss. Formally, our setup can be described as follows:

**The protocol of online multiclass prediction with abstentions**

For each round  $t = 1, 2, \dots, T$ , repeat:

1. The learner observes the  $[K]$ -valued predictions of  $N$  experts, denoted as  $y_{t,1}, \dots, y_{t,N}$ ,
2. based on these observations and possibly some randomness, the learner predicts  $\hat{y}_t \in [K] \cup \{*\}$ ,
3. the environment reveals  $y_t \in [K]$ ,
4. the learner suffers the loss  $\hat{\ell}_t = \mathbb{I}_{\{\hat{y}_t \neq y_t\}}$  if  $\hat{y}_t \in [K]$ , otherwise a loss of  $\hat{\ell}_t = c < \frac{1}{2}$  if  $\hat{y}_t = *$ .

Our algorithm described in the previous sections can be extended to this setup in a straightforward way. Specifically, let us define  $w_{t,i}$  and  $q_{t,i}$  for each expert  $i$  in the same way as before, and define the aggregated probability assigned to class  $k$  as  $p_{t,k} = \sum_{i=1}^N q_{t,i} \mathbb{I}_{\{y_{t,i}=k\}}$ . Given the above notation, Algorithm 1 can be adapted to the multiclass case by making the following adjustments:

1. Redefine  $p_t^* = \max_{k \in [K]} p_{t,k}$  and  $k_t^* = \arg \max_{k \in [K]} p_{t,k}$ , that is, the label attaining the highest probability, and
2. let  $\alpha_t = 2(1 - p_t^*) \wedge 1$  and predict  $*$  with probability  $\alpha_t$  and  $k_t^*$  with probability  $1 - \alpha_t$ .

Note that whenever  $p_t^* \leq \frac{1}{2}$ , the algorithm abstains with probability 1. While this may appear surprising at first sight, it is easily justified by observing that the probability of predicting incorrectly in such cases results in a loss greater than  $\frac{1}{2}$ , thus one can only win by abstaining by our assumption that  $c \leq \frac{1}{2}$ . We also note that the lower bound of Theorem 9 continues to hold in the multiclass setting since one can clearly embed a binary classification problem into this more general problem

setting. Thus, we cannot expect to have rates faster than  $\frac{\log N}{1-2c}$  in the multiclass case either. The following theorem (proved in Appendix B) establishes an upper bound of this exact order, thus matching the lower bound.

**Theorem 12** *Suppose that  $c < \frac{1}{2}$  and  $\eta = 2(1 - 2c) \vee \sqrt{\frac{8 \log N}{T}}$ . Then, the expected regret of our algorithm for multiclass classification satisfies*

$$R_T \leq \frac{\log N}{2(1 - 2c)} \wedge \sqrt{\frac{T \log N}{2}}.$$

## 5.2. Infinite sets of experts

Let us return to the case of binary classification with the twist that we now allow the set of hypotheses to be infinite. As shown by Ben-David, Pál, and Shalev-Shwartz (2009), the complexity of online prediction in this case is characterized by the so-called *Littlestone dimension* of the class  $\mathcal{H}$  of hypothesis—for the precise definition, we refer to Ben-David et al. (2009). Assuming that the Littlestone dimension  $L$  of the class is finite, Ben-David et al. show that a clever extension of the exponentially weighted forecaster satisfies the regret bound  $R_T \leq \sqrt{2LT \log(\frac{eT}{L})}$ .

In our setup we provide the following regret bound, and defer the proof to Appendix B.

**Corollary 13** *Assume that  $\mathcal{H}$  has finite Littlestone dimension  $L$  and  $T \geq L$ , and that the price of abstention is  $c \leq \frac{1}{2}$ . Then, there is a randomized prediction strategy satisfying*

$$R_T \leq \frac{L \log \frac{eT}{L}}{2(1 - 2c)} \wedge \sqrt{\frac{LT \log(\frac{eT}{L})}{2}}.$$

In particular, whenever  $c$  is separated from  $\frac{1}{2}$  the regret scales as  $O(\log T)$  which is much better than the standard  $O(\sqrt{T})$  dependence. We also remark that this result can be seen as an online analog of Theorem 2.2 in Bousquet and Zhivotovskiy (2019), with the Littlestone dimension replacing the VC dimension therein. However, as we mentioned, their algorithm and analysis are based on completely different techniques that take their roots in empirical process theory.

## 6. Concluding remarks and directions of future research

For the first time in the context of online learning, we have shown that endowing the learner with the option of abstaining from prediction allows a significant improvement in the best achievable rates, and that a very natural algorithm can attain the optimal rates. Notably, these results are proved without making any assumption about the sequence of outcomes that the learner has to predict or about the curvature of the loss functions. In this view, our results represent a new flavor of fast rates that, to the best of our knowledge, haven't been explored before in the online learning literature. Arguably, one key factor that makes these fast rates possible is that the learner has access to an action that the comparator strategy cannot use. However, this feature alone cannot explain our results: since we only require the extra action to be marginally better than a uniform random prediction, the learner cannot abuse it to easily beat the best expert. How to characterize properties of such powerful additional actions for the learner in more general convex optimization problems remains to be seen.

Our key technical contribution is observing that the expected loss of our algorithm is upper bounded by the mix loss. This technique played a crucial role in proving a number of classic results in the literature, and was traditionally enabled by directly making an assumption about the curvature of the loss function evaluating the learner’s predictions. To our knowledge, our work is the first to make use of this technique without making direct assumptions about the loss. Rather, our approach is based on transforming the output  $p_t$  of the exponential weights algorithm into a  $\{0, *, 1\}$ -valued prediction in a way that effectively replaces the linear loss by a mixable one from the perspective of the algorithm producing  $p_t$ . Since this transformation can be applied to the output of any algorithm, we find it plausible that our rates could be improved by a more refined prediction algorithm such as the Aggregating Forecaster of [Vovk \(1998, 2001\)](#). Another possible improvement to pursue is proving high-probability versions of our upper bounds.

Another important technique we employed for proving our bounds in Section 3 was an adaptive stepsize schedule that decreases the learning rate in each round when it proves too large for the loss to be mixable. This technique bears a vague resemblance to the AdaHedge algorithm of [Van Erven et al. \(2011\)](#); [De Rooij et al. \(2014\)](#) that uses a learning rate proportional to the cumulative gap between the expected loss and the mix loss. However, their analysis is restricted to the case where this gap is nonnegative, whereas our bounds crucially use that they can be often negative. As a result, our technique can be directly adapted to account for more general mixable losses with time-dependent mixability parameters. Similarly, our variant of Tsybakov’s condition on the abstention costs can be adapted to mixability parameters to characterize the hardness of learning with a non-stationary sequence of loss functions. For instance, one can easily prove the following result by adjusting our proof techniques in a straightforward way:

**Proposition 14** *Consider the problem of predicting a sequence of outcomes in  $[0, 1]$  under a sequence of loss functions  $\ell_{q_t} = |\cdot|^{q_t}$ ,  $q_t \in (1, 2]$  that satisfy the following Tsybakov’s type margin assumption for  $q_t$ : for any  $x > 0$ ,  $\frac{1}{T} \sum_{t=1}^T \mathbb{I}_{\{q_t-1 < x\}} \leq \beta x^{\frac{\alpha}{1-\alpha}}$  holds for some  $\beta > 0, \alpha \in [0, 1)$ . Then, there is a deterministic prediction strategy satisfying  $R_T = O\left((\log N)^{\frac{1}{2-\alpha}} T^{\frac{1-\alpha}{2-\alpha}}\right)$ .*

This result uses a variant of Tsybakov’s condition that characterizes the behaviour of the second derivative of the loss around zero. The statement is proven by observing that the expected loss is upper-bounded by the mix loss whenever  $\eta \leq q - 1$ , and then closely following the proof steps in Section 3—we leave this as an easy exercise for the reader. The result can be seen to address an open question of [Haussler et al. \(1998\)](#) about whether intermediate rates between  $O(\log N)$  and  $O(\sqrt{T \log N})$  are attainable in online prediction under suitable curvature assumptions on the loss functions. It also bears resemblance to the results of [Bartlett, Hazan, and Rakhlin \(2008\)](#) who provide algorithms that adapt to the curvature of the losses in more general online convex optimization problems, the key difference being that the curvature of our losses can be quantified in terms of exp-concavity rather than strong convexity as done by [Bartlett et al.](#)

Finally, we note that while our techniques are different from those of [Bousquet and Zhivotovskiy \(2019\)](#) who prove analogous results for the setting of statistical learning, the two analyses actually related in a very subtle way: both of them are based on techniques that are traditionally used for proving fast rates for losses with curvature, in their respective settings. While our approach is based on bounding the mixability gap as discussed above, the one of [Bousquet and Zhivotovskiy](#) is based on an extension of the recently proposed technique of [Mendelson \(2019\)](#) for statistical learning with the square loss. Our lower bound also makes use of this connection. Deeper investigation of this connection between curvature and the abstention may lead to further interesting results.

## Acknowledgments

We thank the three anonymous reviewers for their valuable feedback that helped us improve the paper, particularly for catching a small bug in the original proof of Theorem 12. G. Neu was supported by “la Caixa” Banking Foundation through the Junior Leader Postdoctoral Fellowship Programme, a Google Faculty Research Award, and a Bosch AI Young Researcher Award.

## References

- J.-Y. Audibert. Fast learning rates in statistical inference through aggregation. *The Annals of Statistics*, 37(4):1591–1646, 2009.
- P. L. Bartlett and S. Mendelson. Empirical minimization. *Probability theory and related fields*, 135(3):311–334, 2006.
- P. L. Bartlett and M. H. Wegkamp. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9(Aug):1823–1840, 2008.
- P. L. Bartlett, E. Hazan, and A. Rakhlin. Adaptive online gradient descent. In *Advances in Neural Information Processing Systems*, pages 65–72, 2008.
- S. Ben-David, D. Pál, and S. Shalev-Shwartz. Agnostic online learning. In *COLT*, volume 3, page 1, 2009.
- O. Bousquet and N. Zhivotovskiy. Fast classification rates without standard margin assumptions. *arXiv preprint arXiv:1910.12756*, 2019.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- C. Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1):41–46, 1970.
- C. Cortes, G. DeSalvo, C. Gentile, M. Mohri, and S. Yang. Online learning with abstention. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1059–1067. PMLR, 2018.
- T. M. Cover. Behavior of sequential predictors of binary sequences. Technical report, Trans. 4th Prague Conf. Information Theory Statistical Decision Functions, Random Processes, 1965.
- S. De Rooij, T. Van Erven, P. D. Grünwald, and W. M. Koolen. Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*, 15(1):1281–1316, 2014.
- Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55:119–139, 1997.
- J. Hannan. Approximation to Bayes risk in repeated play. *Contributions to the theory of games*, 3: 97–139, 1957.
- D. Haussler, J. Kivinen, and M. K. Warmuth. Sequential prediction of individual sequences under general loss functions. *IEEE Transactions on Information Theory*, 44(5):1906–1925, 1998.

- R. Herbei and M. H. Wegkamp. Classification with reject option. *Canadian Journal of Statistics*, 34(4):709–721, 2006.
- W. M. Koolen, P. Grünwald, and T. van Erven. Combining adversarial guarantees and stochastic fast rates in online learning. In *Advances in Neural Information Processing Systems*, pages 4457–4465, 2016.
- N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.
- E. Mammen and A. B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- P. Massart and É. Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- S. Mendelson. An unrestricted learning procedure. *J. ACM*, 66(6), 2019.
- O. Shamir. The sample complexity of learning linear predictors with the squared loss. *The Journal of Machine Learning Research*, 16(1):3475–3486, 2015.
- A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- T. Van Erven, W. M. Koolen, S. D. Rooij, and P. Grünwald. Adaptive hedge. In *Advances in Neural Information Processing Systems*, pages 1656–1664, 2011.
- T. van Erven, P. D. Grünwald, N. A. Mehta, M. D. Reid, and R. C. Williamson. Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16(54):1793–1861, 2015.
- V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Proc. USSR Acad. Sci.*, 1968.
- V. Vovk. Aggregating strategies. *Proc. of Computational Learning Theory*, 1990.
- V. Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56(2):153–173, 1998.
- V. Vovk. Competitive on-line statistics. *International Statistical Review*, 69(2):213–248, 2001.
- C. Zhang and K. Chaudhuri. The extended Littlestone’s dimension for learning with mistakes and abstentions. In *Conference on Learning Theory*, pages 1584–1616, 2016.

## Appendix A. Why do we need to randomize?

**Example 1** Consider the case of two experts such that one of them always predicts zero and the other always predicts one. Assume that the environment waits for the learner’s prediction and then provides the opposite label as the true label if the learner predicts zero or one. If the learner decides to abstain the environment provides the label according to the prediction of the best expert on  $T$  rounds. In what follows, we show that the regret of any deterministic strategy with abstention will be at least  $\frac{cT}{2}$ .

Assume that out of  $T$  rounds the learner used the abstention  $k$  times and on the remaining  $T - k$  rounds they predicted 0 or 1. The total loss suffered by the learner will be equal to  $ck + T - k$ . Now we need to upper bound the loss of the best expert. The simple upper bound is  $\frac{T}{2}$ . Now, since the environment is aware of the deterministic strategy of the learner we can always guarantee that it reveals the labels such that the loss of the best expert on these  $k$  rounds is equal to zero. This implies the number of mistakes made by the best expert is at most  $\frac{T}{2} \wedge (T - k)$ . Considering the case  $k \geq \frac{T}{2}$  we have that the regret is at least  $ck \geq \frac{cT}{2}$ . If  $k < \frac{T}{2}$  we have that the regret is at least  $(\frac{T}{2} - 1)(c - 1) + \frac{T}{2} \geq \frac{cT}{2} + \frac{1}{2}$ .

## Appendix B. Omitted proofs

### Proof of Lemma 3

The proof is based on studying the evolution of the cumulative weights  $W_t = \sum_i w_{t,i}$ . Defining  $w_{1,i} = 1$  and  $L_{T,i} = \sum_{k=1}^T \ell_{t,i}$  for all  $i$ , we have

$$\begin{aligned} -\min_i L_{T,i} &= \min_i \frac{1}{\eta} \log e^{-\eta L_{T,i}} \leq \frac{1}{\eta} \log \sum_i e^{-\eta L_{T,i}} = \frac{1}{\eta} \log W_{T+1} = \frac{1}{\eta} \sum_{t=1}^T \log \frac{W_{t+1}}{W_t} + \frac{\log W_1}{\eta} \\ &= \frac{1}{\eta} \sum_{t=1}^T \log \frac{\sum_i w_{t,i} e^{-\eta \ell_{t,i}}}{\sum_j w_{t,j}} + \frac{\log N}{\eta} = \frac{1}{\eta} \sum_{t=1}^T \log \sum_i q_{t,i} e^{-\eta \ell_{t,i}} + \frac{\log N}{\eta}. \end{aligned}$$

Reordering gives the result. ■

### Proof of Proposition 5

The proof is based on considering two types of rounds depending on the value of  $c_t$ . In the case where  $c_t$  is such that  $\eta \leq 2(1 - 2c_t)$  is satisfied, we have  $\mathbb{E}[\widehat{\ell}_t] \leq \widetilde{\ell}_t$  by Lemma 4. Otherwise, we use the general bound  $\mathbb{E}[\widehat{\ell}_t] \leq \widetilde{\ell}_t + \frac{\eta}{8}$  as in the proof of Corollary 2. The proof is concluded by appealing to Lemma 3. ■

### Proof of Theorem 8

We begin by studying some properties of the learning rate  $\eta^*$  optimizing the function  $B_T(\eta) = \frac{\log N}{\eta} + \frac{\eta \tau(\eta)}{8}$ , where  $\tau(\eta) = \sum_{t=1}^T \mathbb{I}_{\{2(1-2c_t) \leq \eta\}}$ . We observe that  $B_T$  is lower-semicontinuous with at most  $T$  discontinuities, and thus it takes its minimum either at one of its discontinuities or when its derivative is zero. In either case,  $\eta^*$  satisfies

$$\frac{\log N}{\eta^*} \geq \frac{\eta^* \tau(\eta^*)}{8},$$

with equality when the minimum is achieved where the function is continuous. Note that this implies that  $\tau^* = \tau(\eta^*)$  is bounded as  $\tau^* \leq 8 \log N / (\eta^*)^2$ . Furthermore, the bound also entails  $R_T^* \geq \frac{2 \log N}{\eta^*}$ .

Regarding the regret of our algorithm, standard arguments can be used to prove the bound

$$R_T \leq \frac{\log N}{\eta_{T+1}} + \sum_{t=1}^T \left( \mathbb{E}[\widehat{\ell}_t] - \widetilde{\ell}_t \right),$$

which actually holds for general nonincreasing learning-rate sequences (see, e.g., Lemma 2 in De Rooij et al., 2014). By the same argument as used to prove Corollary 2, each term in the sum on the right-hand side of the above bound can be bounded as

$$\mathbb{E}[\widehat{\ell}_t] - \widetilde{\ell}_t \leq \frac{\eta_t}{8} \mathbb{I}_{\{2(1-2c_t) < \eta_t\}},$$

so that the sum itself can be bounded as

$$\sum_{t=1}^T \left( \mathbb{E}[\widehat{\ell}_t] - \widetilde{\ell}_t \right) \leq \sum_{t=1}^T \frac{\eta_t}{8} \mathbb{I}_{\{2(1-2c_t) < \eta_t\}} \leq \frac{1}{8} \sum_{k=1}^{d_T} \sqrt{\frac{\log N}{k}} \leq \frac{1}{4} \sqrt{d_T \log N},$$

where we used that  $d_t$  increases precisely whenever  $\mathbb{I}_{\{2(1-2c_t) < \eta_t\}} = 1$ , and that  $\sum_{k=1}^K \sqrt{1/k} \leq 2\sqrt{K}$ . Now notice that by monotonicity of  $\tau(\eta)$ , we have  $d_T > \tau^*$ , which implies that there exists a  $t^* \leq T$  such that  $d_t > d^* = \log N / (\eta^*)^2$  holds for all  $t > t^*$ , so that we have

$$\eta_t = \sqrt{\frac{\log N}{d_t}} \leq \sqrt{\frac{\log N}{d^*}} = \eta^*.$$

After round  $t^*$ , the remaining number of times that our algorithm updates the learning rate is  $\sum_{t=t^*+1}^T \mathbb{I}_{\{2(1-2c_t) \leq \eta_t\}}$ , which can be bounded as

$$\sum_{t=t^*+1}^T \mathbb{I}_{\{2(1-2c_t) \leq \eta_t\}} \leq \sum_{t=t^*+1}^T \mathbb{I}_{\{2(1-2c_t) \leq \eta^*\}} \leq \sum_{t=1}^T \mathbb{I}_{\{2(1-2c_t) \leq \eta^*\}} = \tau^*,$$

where the first inequality uses the monotonicity of each of the summands in  $\eta$  and the second step adds a number of non-negative terms to the sum. The final step is the definition of  $\tau^*$ . Altogether, this implies the bound  $d_T \leq d^* + 1 + \tau^* \leq 9 \log N / (\eta^*)^2 + 1$ , so the cumulative mixability gap is bounded as

$$\sum_{t=1}^T \left( \mathbb{E}[\widehat{\ell}_t] - \widetilde{\ell}_t \right) \leq \frac{1}{4} \sqrt{d_T \log N} \leq \frac{3 \log N}{4\eta^*} + \frac{\sqrt{\log N}}{4}.$$

On the other hand, we have

$$\frac{\log N}{\eta_{T+1}} = \sqrt{d_T \log N} \leq \frac{3 \log N}{\eta^*} + \sqrt{\log N}.$$

Thus, the total regret of our algorithm is bounded as

$$\frac{15 \log N}{4\eta^*} + \frac{5\sqrt{\log N}}{4}.$$

Comparing this bound with the lower bound on  $R_T^*$  concludes the proof.  $\blacksquare$

### Proof of Theorem 9

Our trick will be to construct a randomized prediction algorithm in the batch setup using an algorithm  $\mathcal{A}$  for our main setting through the online-to-batch conversion described in Lemma 10. First,



we fix  $B = \frac{1}{2}$  and observe that it follows from the proof of Lemma 11 that the class of functions  $G$  used in the lower bound is a cube of  $\{\pm\frac{1}{2}\}$ -valued functions consisting of  $\lfloor \log_2 N \rfloor$  unique functions and that we are only interested in distributions with  $Y \in \{\pm\frac{1}{2}\}$ . For the exact details implying these two facts we refer to the definition of the hypercube of probability distributions (Definition 8.1 in Audibert, 2009) and the first lines of the proof of Theorem 8.3 and 8.4 where Audibert sets  $h_1 = B, h_2 = -B$ . Finally, his proof uses  $p_+ = 1, p_- = 0$  and this implies that the  $\min_g \mathbb{E} [|g(X) - Y|^q]$  among all measurable functions is achieved by some  $\{\pm\frac{1}{2}\}$ -valued function for each distribution in the hypercube of  $\lfloor \log_2 N \rfloor$  distributions.

With these important observations in mind we will construct the following game of prediction with expert advice. For each round  $t = 1, 2, \dots, T$ , the environment draws  $(X_t, Y_t) \sim P$ , sets  $y_t = Y_t + \frac{1}{2}$ , and each expert  $i \in [N]$  (we can naturally identify  $\{g_1, g_2, \dots, g_N\}$  with experts) predicts  $y_{t,i} = g_i(X_t) + \frac{1}{2}$ .

This setup allows us to use our online learning algorithm  $\mathcal{A}$  to produce a randomized hypothesis  $\hat{g}$  defined on  $\mathcal{X}$  and taking its values in  $\{\pm\frac{1}{2}, 0\}$  and matching the bound of Lemma 10. For each  $t = 1, 2, \dots, T$  matching the conditions of Lemma 10 we run  $\mathcal{A}$  on  $(X_1, Y_1), \dots, (X_{t-1}, Y_{t-1})$ . For any  $x \in \mathcal{X}$ ,  $\hat{g}$  produces its prediction by feeding the expert predictions  $y'_{t,i} = g_i(x)$  to  $\mathcal{A}$ , obtaining a random prediction  $\hat{y}_T \in \{0, *, 1\}$ , and finally mapping  $\hat{y}_T$  to an element of  $\mathcal{Y}$  according to  $0 \mapsto -\frac{1}{2}$ ,  $1 \mapsto \frac{1}{2}$  and  $*$   $\mapsto 0$ . Observe that since  $Y \in \{\pm\frac{1}{2}\}$  the following holds if we fix  $q = \log_2 \frac{1}{c}$ ,

$$|Y - \hat{g}(x)|^q = \begin{cases} \mathbb{I}_{\{\hat{g}(x) \neq Y\}}, & \text{if } \hat{g}(x) = \pm\frac{1}{2}; \\ c, & \text{if } \hat{g}(x) = 0. \end{cases}$$

At the same time, for any  $g \in G$ , it holds that  $|Y - \hat{g}(x)|^q = \mathbb{I}_{\{\hat{g}(x) \neq Y\}}$ . To conclude, we have in this setup that the loss of  $\hat{g}$  is binary whenever the output is  $\pm\frac{1}{2}$  or equal to  $c$  whenever the output is equal to zero regardless of the value of  $Y \in \{\pm\frac{1}{2}\}$ . Using Lemma 10 we have

$$\mathbb{E} [|Y - \hat{g}(X)|^q] = \frac{1}{T} \mathbb{E} \left[ \sum_{t=1}^T \hat{\ell}_t \right] \quad \text{and} \quad \min_{g \in G} \mathbb{E} [|Y - g(X)|^q] = \min_{g \in G} \mathbb{E} [\mathbb{I}_{\{g(X) \neq Y\}}]. \quad (5)$$

Now, provided that  $c \geq \frac{1}{4}$  for our choice  $q = \log_2 \frac{1}{c}$  we have  $q \leq 2$  and

$$\frac{1}{2} - \frac{1}{2}(q-1) \log 2 \leq \frac{1}{2^q} \leq \frac{1}{2} - \frac{1}{4}(q-1) \log 2,$$

which implies

$$\frac{1}{4}(q-1) \log 2 \leq \frac{1}{2} - c \leq \frac{1}{2}(q-1) \log 2. \quad (6)$$

Since  $T \geq 4 \log N$  in order to apply Lemma 11 we need  $q \geq 1 + \sqrt{\frac{\lfloor \log_2 N \rfloor}{4T}}$  which holds if

$$\frac{1}{2} - c \geq \log 2 \sqrt{\frac{\lfloor \log_2 N \rfloor}{4T}}. \quad (7)$$

This corresponds to the regime where the abstentions can give some gain compared to the slow rate bound  $O(\sqrt{T \log N})$ . Therefore, provided that  $c \geq \frac{1}{4}$  since the lower bound of Lemma 11 holds for any  $[-\frac{1}{2}, \frac{1}{2}]$ -valued estimator (and therefore, for  $\{\pm\frac{1}{2}, 0\}$ -valued  $\hat{g}$ ) we have for some  $C > 0$ ,

$$\mathbb{E} \left[ \sum_{t=1}^T \hat{\ell}_t \right] \geq T \min_{g \in G} \mathbb{E} [\mathbb{I}_{\{Y \neq g(X)\}}] + C \frac{\log N}{1 - 2c},$$

where we used (5). We remark that the bound in Lemma 11 has the factor  $\frac{q}{90(q-1)} \vee e^{-1}$  with the first term being dominant if  $1 \leq q \leq 1.031$ . However, it can be easily seen that the values of  $q > 1.031$  can be controlled by the choice of the absolute constant  $C$  above. Finally, by Jensen's inequality

$$T \min_{g \in G} \mathbb{E} [\mathbb{I}_{\{Y \neq g(X)\}}] = \min_{g \in G} \sum_{t=1}^T \mathbb{E} [\mathbb{I}_{\{Y_t \neq g(X_t)\}}] \geq \mathbb{E} \left[ \min_{g \in G} \sum_{t=1}^T \mathbb{I}_{\{Y_t \neq g(X_t)\}} \right] = \mathbb{E} \left[ \min_i \sum_{t=1}^T \ell_{t,i} \right].$$

This implies that  $T \min_{g \in G} \mathbb{E} [\mathbb{I}_{\{Y \neq g(X)\}}]$  is greater than the expected regret of the best expert. Therefore, we show that for any randomized algorithm with abstentions there is a random environment such that the expected regret is  $\Omega\left(\frac{\log N}{1-2c}\right)$  provided that  $q = \log_2 \frac{1}{c}$  satisfies  $q - 1 \geq \sqrt{\frac{\lfloor \log_2 N \rfloor}{4T}}$ . This, of course, implies the existence of the deterministic strategy for the environment.

To finish the proof we need to consider the regime

$$1 \leq q \leq 1 + \sqrt{\frac{\lfloor \log_2 N \rfloor}{4T}}.$$

and show that it gives the lower bound of order  $\Omega(\sqrt{T \log N})$  which does not depend of  $c$  anymore. Since the total loss of any algorithm grows monotonically as  $c$  approaches  $\frac{1}{2}$ , we can prove the lower bound only for  $q = 1 + \sqrt{\frac{\lfloor \log_2 N \rfloor}{4T}}$ , where as before  $q = \log_2 \frac{1}{c}$ . Indeed, the lower bound will extend automatically for smaller values of  $q$  (this corresponds to larger values of  $c$ ). Fortunately, that value of  $q$  is covered by the first part of the proof. It is left to observe that in this case we have that  $\frac{\log N}{1-2c}$  is controlled by  $\sqrt{T \log N}$  up to a multiplicative constant factor. The claim follows.  $\blacksquare$

### Proof of Theorem 12

We begin by defining the misclassification probability

$$r_t = \sum_{k=1}^K p_{t,k} \mathbb{I}_{\{k \neq y_t\}} = \sum_{i=1}^N q_{t,i} \mathbb{I}_{\{y_{t,i} \neq y_t\}}$$

and noticing that the mix loss can be again written as

$$-\frac{1}{\eta} \log \sum_i q_{t,i} e^{-\eta \ell_{t,i}} = -\frac{1}{\eta} \log (1 + r_t (e^{-\eta} - 1)).$$

On the other hand, the expected loss of our algorithm is

$$\widehat{\ell}_t = \alpha_t c + (1 - \alpha_t) \mathbb{I}_{\{k_t^* \neq y_t\}}.$$

Notice that this equals  $c$  when  $p_t^* < \frac{1}{2}$ . In this case, the error probability satisfies  $r_t \geq \frac{1}{2}$  so that the mix loss is can be lower bounded as

$$\begin{aligned} -\frac{1}{\eta} \log (1 + r_t (e^{-\eta} - 1)) &\geq -\frac{1}{\eta} \log \left( \frac{1}{2} (1 + e^{-\eta}) \right) \geq -\frac{1}{\eta} \log \left( 1 - \frac{\eta}{2} + \frac{\eta^2}{4} \right) \geq \frac{1}{2} - \frac{\eta}{4} \\ &\geq \frac{1}{2} - \frac{2(1-2c)}{4} = c, \end{aligned}$$

where in the first line we used the inequalities  $e^{-x} \leq 1 - x + \frac{x^2}{2}$  and  $\log(1 - x) \leq -x$  that both hold for  $x \geq 0$ , and our assumption that  $\eta \leq 2(1 - 2c)$  in the last line. Thus, the expected loss lower bounds the mix loss in this case.

Considering the case where  $p_t^* \geq \frac{1}{2}$ , we have  $r_t = 1 - p_t^* \leq \frac{1}{2}$  whenever  $k_t^* = y_t$  holds, and  $\frac{1}{2} \leq p_t^* \leq r_t$  otherwise (since in this case we have  $r_t = \sum_{k \neq y_t} p_{k,t} \geq p_t^*$ ). In the first case, we have

$$\mathbb{E}[\widehat{\ell}_t] = \alpha_t c = 2(1 - p_t^*)c = 2r_t c = r_t + (2c - 1)r_t,$$

while in the second one we have

$$\mathbb{E}[\widehat{\ell}_t] = \alpha_t(c - 1) + 1 = 2(p_t^* - 1)(1 - c) + 1 \leq 2(r_t - 1)(1 - c) + 1 = r_t + (2c - 1)(1 - r_t).$$

Summarizing both cases, we have

$$\mathbb{E}[\widehat{\ell}_t] \leq r_t + (2c - 1)(r_t \wedge (1 - r_t)).$$

The proof is concluded by noticing that this function is upper-bounded by the mix loss  $\widetilde{\ell}_t$  by the same argument as used in the proof of Lemma 4, and using Lemma 3 to upper-bound the cumulative mix loss.  $\blacksquare$

### Proof of Corollary 13

Our key technical tool is the following lemma which can be seen as an online learning analog of the renowned lemma by Vapnik and Chervonenkis (1968).

**Lemma 15 (Ben-David et al. (2009))** *Fix  $T$  and assume that  $\mathcal{H}$  has the Littlestone dimension  $L$ . There is a set  $\mathcal{H}_T$  of at most  $\sum_{i=0}^L \binom{T}{i}$  experts such that for any  $h \in \mathcal{H}$  there is at least one expert in  $\mathcal{H}_T$  predicting as  $h$  on any sequence of length at most  $T$ .*

Now the proof follows immediately from Corollary 2, Lemma 15 and  $\sum_{i=0}^L \binom{T}{i} \leq \left(\frac{eT}{L}\right)^L$  whenever  $T \geq L$ . We remark that the results of Section 3 can also be straightforwardly adapted for the sets of experts having finite Littlestone dimension.  $\blacksquare$