
Mirror Descent Meets Fixed Share (and feels no regret)

Nicolò Cesa-Bianchi
Università degli Studi di Milano
nicolo.cesa-bianchi@unimi.it

Pierre Gaillard
Ecole Normale Supérieure,* Paris
pierre.gaillard@ens.fr

Gábor Lugosi
ICREA & Universitat Pompeu Fabra, Barcelona
gabor.lugosi@upf.edu

Gilles Stoltz
Ecole Normale Supérieure,* Paris &
HEC Paris, Jouy-en-Josas, France
gilles.stoltz@ens.fr

Abstract

Mirror descent with an entropic regularizer is known to achieve shifting regret bounds that are logarithmic in the dimension. This is done using either a carefully designed projection or by a weight sharing technique. Via a novel unified analysis, we show that these two approaches deliver essentially equivalent bounds on a notion of regret generalizing shifting, adaptive, discounted, and other related regrets. Our analysis also captures and extends the generalized weight sharing technique of Bousquet and Warmuth, and can be refined in several ways, including improvements for small losses and adaptive tuning of parameters.

1 Introduction

Online convex optimization is a sequential prediction paradigm in which, at each time step, the learner chooses an element from a fixed convex set \mathcal{S} and then is given access to a convex loss function defined on the same set. The value of the function on the chosen element is the learner's loss. Many problems such as prediction with expert advice, sequential investment, and online regression/classification can be viewed as special cases of this general framework. Online learning algorithms are designed to minimize the regret. The standard notion of regret is the difference between the learner's cumulative loss and the cumulative loss of the single best element in \mathcal{S} . A much harder criterion to minimize is shifting regret, which is defined as the difference between the learner's cumulative loss and the cumulative loss of an arbitrary sequence of elements in \mathcal{S} . Shifting regret bounds are typically expressed in terms of the *shift*, a notion of regularity measuring the length of the trajectory in \mathcal{S} described by the comparison sequence (i.e., the sequence of elements against which the regret is evaluated). In online convex optimization, shifting regret bounds for convex subsets $\mathcal{S} \subseteq \mathbb{R}^d$ are obtained for the projected online mirror descent (or follow-the-regularized-leader) algorithm. In this case the shift is typically computed in terms of the p -norm of the difference of consecutive elements in the comparison sequence —see [1, 2] and [3].

We focus on the important special case when \mathcal{S} is the simplex. In [1] shifting bounds are shown for projected mirror descent with entropic regularizers using a 1-norm to measure the shift.¹ When the comparison sequence is restricted to the corners of the simplex (which is the setting of prediction with expert advice), then the shift is naturally defined to be the number of times the trajectory moves

*Ecole Normale Supérieure, Paris – CNRS – INRIA, within the project-team CLASSIC

¹Similar 1-norm shifting bounds can also be proven using the analysis of [2]. However, without using entropic regularizers it is not clear how to achieve a logarithmic dependence on the dimension, which is one of the advantages of working in the simplex.

to a different corner. This problem is often called “tracking the best expert” —see, e.g., [4, 5, 1, 6, 7], and it is well known that exponential weights with weight sharing, which corresponds to the fixed-share algorithm of [4], achieves a good shifting bound in this setting. In [6] the authors introduce a generalization of the fixed-share algorithm, and prove various shifting bounds for any trajectory in the simplex. However, their bounds are expressed using a quantity that corresponds to a proper shift only for trajectories on the simplex corners.

In this paper we offer a unified analysis of mirror descent, fixed share, and the generalized fixed share of [6] for the setting of online convex optimization in the simplex. Our bounds are expressed in terms of a notion of shift based on the total variation distance. Our analysis relies on a generalized notion of shifting regret which includes, as special cases, related notions of regret such as adaptive regret, discounted regret, and regret with time-selection functions. Perhaps surprisingly, we show that projected mirror descent and fixed share achieve essentially the same generalized regret bound. Finally, we show that widespread techniques in online learning, such as improvements for small losses and adaptive tuning of parameters, are all easily captured by our analysis.

2 Preliminaries

For simplicity, we derive our results in the setting of online linear optimization. As we show in the supplementary material, these results can be easily extended to the more general setting of online convex optimization through a standard linearization step.

Online linear optimization may be cast as a repeated game between the *forecaster* and the *environment* as follows. We use Δ_d to denote the simplex $\{\mathbf{q} \in [0, 1]^d : \|\mathbf{q}\|_1 = 1\}$.

Online linear optimization in the simplex. For each round $t = 1, \dots, T$,

1. Forecaster chooses $\hat{\mathbf{p}}_t = (\hat{p}_{1,t}, \dots, \hat{p}_{d,t}) \in \Delta_d$
2. Environment chooses a loss vector $\boldsymbol{\ell}_t = (\ell_{1,t}, \dots, \ell_{d,t}) \in [0, 1]^d$
3. Forecaster suffers loss $\hat{\mathbf{p}}_t^\top \boldsymbol{\ell}_t$.

The goal of the forecaster is to minimize the accumulated loss, e.g., $\widehat{L}_T = \sum_{t=1}^T \hat{\mathbf{p}}_t^\top \boldsymbol{\ell}_t$. In the now classical problem of prediction with expert advice, the goal of the forecaster is to compete with the best fixed component (often called “expert”) chosen in hindsight, that is, with $\min_{i=1, \dots, T} \sum_{t=1}^T \ell_{i,t}$; or even to compete with a richer class of *sequences* of components. In Section 3 we state more specifically the goals considered in this paper.

We start by introducing our main algorithmic tool, described in Figure 1, a share algorithm whose formulation generalizes the seemingly unrelated formulations of the algorithms studied in [4, 1, 6]. It is parameterized by the “mixing functions” $\psi_t : [0, 1]^d \rightarrow \Delta_d$ for $t \geq 2$ that assign probabilities to past “pre-weights” as defined below. In all examples discussed in this paper, these mixing functions are quite simple, but working with such a general model makes the main ideas more transparent. We then provide a simple lemma that serves as the starting point² for analyzing different instances of this generalized share algorithm.

Lemma 1. *For all $t \geq 1$ and for all $\mathbf{q}_t \in \Delta_d$, Algorithm 1 satisfies*

$$(\hat{\mathbf{p}}_t - \mathbf{q}_t)^\top \boldsymbol{\ell}_t \leq \frac{1}{\eta} \sum_{i=1}^d q_{i,t} \ln \frac{v_{i,t+1}}{\hat{p}_{i,t}} + \frac{\eta}{8}.$$

Proof. By Hoeffding’s inequality (see, e.g., [3, Section A.1.1]),

$$\sum_{j=1}^d \hat{p}_{j,t} \ell_{j,t} \leq -\frac{1}{\eta} \ln \left(\sum_{j=1}^d \hat{p}_{j,t} e^{-\eta \ell_{j,t}} \right) + \frac{\eta}{8}. \quad (1)$$

By definition of $v_{i,t+1}$, for all $i = 1, \dots, d$ we then have $\sum_{j=1}^d \hat{p}_{j,t} e^{-\eta \ell_{j,t}} = \hat{p}_{i,t} e^{-\eta \ell_{i,t}} / v_{i,t+1}$, which implies $\hat{\mathbf{p}}_t^\top \boldsymbol{\ell}_t \leq \ell_{i,t} + (1/\eta) \ln(v_{i,t+1} / \hat{p}_{i,t}) + \eta/8$. The proof is concluded by taking a convex aggregation with respect to \mathbf{q}_t . \square

²We only deal with linear losses in this paper. However, it is straightforward that for sequences of η -exp-concave loss functions, the additional term $\eta/8$ in the bound is no longer needed.

Parameters: learning rate $\eta > 0$ and mixing functions ψ_t for $t \geq 2$

Initialization: $\hat{\mathbf{p}}_1 = \mathbf{v}_1 = (1/d, \dots, 1/d)$

For each round $t = 1, \dots, T$,

1. Predict $\hat{\mathbf{p}}_t$;

2. Observe loss $\ell_t \in [0, 1]^d$;

3. [loss update] **For each** $j = 1, \dots, d$ define

$$v_{j,t+1} = \frac{\hat{p}_{j,t} e^{-\eta \ell_{j,t}}}{\sum_{i=1}^d \hat{p}_{i,t} e^{-\eta \ell_{i,t}}} \quad \text{the current pre-weights, and } \mathbf{v}_{t+1} = (v_{1,t+1}, \dots, v_{d,t+1});$$

$$\underline{V}_{t+1} = [v_{i,s}]_{1 \leq i \leq d, 1 \leq s \leq t+1} \quad \text{the } d \times (t+1) \text{ matrix of all past and current pre-weights;}$$

4. [shared update] Define $\hat{\mathbf{p}}_{t+1} = \psi_{t+1}(\underline{V}_{t+1})$.

Algorithm 1: The generalized share algorithm.

3 A generalized shifting regret for the simplex

We now introduce a generalized notion of shifting regret which unifies and generalizes the notions of discounted regret (see [3, Section 2.11]), adaptive regret (see [8]), and shifting regret (see [2]). For a fixed horizon T , a sequence of discount factors $\beta_{t,T} \geq 0$ for $t = 1, \dots, T$ assigns varying weights to the instantaneous losses suffered at each round. We compare the total loss of the forecaster with the loss of an arbitrary sequence of vectors $\mathbf{q}_1, \dots, \mathbf{q}_T$ in the simplex Δ_d . Our goal is to bound the regret

$$\sum_{t=1}^T \beta_{t,T} \hat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \beta_{t,T} \mathbf{q}_t^\top \ell_t$$

in terms of the “regularity” of the comparison sequence $\mathbf{q}_1, \dots, \mathbf{q}_T$ and of the variations of the discounting weights $\beta_{t,T}$. By setting $\mathbf{u}_t = \beta_{t,T} \mathbf{q}_t^\top \in \mathbb{R}_+^d$, we can rephrase the above regret as

$$\sum_{t=1}^T \|\mathbf{u}_t\|_1 \hat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \mathbf{u}_t^\top \ell_t. \quad (2)$$

In the literature on tracking the best expert [4, 5, 1, 6], the regularity of the sequence $\mathbf{u}_1, \dots, \mathbf{u}_T$ is measured as the number of times $\mathbf{u}_t \neq \mathbf{u}_{t+1}$. We introduce the following regularity measure

$$m(\mathbf{u}_1^T) = \sum_{t=2}^T D_{\text{TV}}(\mathbf{u}_t, \mathbf{u}_{t-1}) \quad (3)$$

where for $\mathbf{x} = (x_1, \dots, x_d), \mathbf{y} = (y_1, \dots, y_d) \in \mathbb{R}_+^d$, we define $D_{\text{TV}}(\mathbf{x}, \mathbf{y}) = \sum_{x_i \geq y_i} (x_i - y_i)$. Note that when $\mathbf{x}, \mathbf{y} \in \Delta_d$, we recover the total variation distance $D_{\text{TV}}(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_1$, while for general $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^d$, the quantity $D_{\text{TV}}(\mathbf{x}, \mathbf{y})$ is not necessarily symmetric and is always bounded by $\|\mathbf{x} - \mathbf{y}\|_1$. The traditional shifting regret of [4, 5, 1, 6] is obtained from (2) when all \mathbf{u}_t are such that $\|\mathbf{u}_t\|_1 = 1$.

4 Projected update

The shifting variant of the EG algorithm analyzed in [1] is a special case of the generalized share algorithm in which the function ψ_{t+1} performs a projection of the pre-weights on the convex set $\Delta_d^\alpha = [\alpha/d, 1]^d \cap \Delta_d$. Here $\alpha \in (0, 1)$ is a fixed parameter. We can prove (using techniques similar to the ones shown in the next section—see the supplementary material) the following bound which generalizes [1, Theorem 16].

Theorem 1. For all $T \geq 1$, for all sequences $\ell_1, \dots, \ell_T \in [0, 1]^d$ of loss vectors, and for all $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathbb{R}_+^d$, if Algorithm 1 is run with the above update, then

$$\sum_{t=1}^T \|\mathbf{u}_t\|_1 \widehat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \mathbf{u}_t^\top \ell_t \leq \frac{\|\mathbf{u}_1\|_1 \ln d}{\eta} + \frac{m(\mathbf{u}_1^T)}{\eta} \ln \frac{d}{\alpha} + \left(\frac{\eta}{8} + \alpha\right) \sum_{t=1}^T \|\mathbf{u}_t\|_1. \quad (4)$$

This bound can be optimized by a proper tuning of α and η parameters. We show a similarly tuned (and slightly better) bound in Corollary 1.

5 Fixed-share update

Next, we consider a different instance of the generalized share algorithm corresponding to the update

$$\widehat{p}_{j,t+1} = \sum_{i=1}^d \left(\frac{\alpha}{d} + (1-\alpha)\mathbb{1}_{i=j} \right) v_{i,t+1} = \frac{\alpha}{d} + (1-\alpha)v_{j,t+1}, \quad 0 \leq \alpha \leq 1 \quad (5)$$

Despite seemingly different statements, this update in Algorithm 1 can be seen to lead *exactly* to the fixed-share algorithm of [4] for prediction with expert advice. We now show that this update delivers a bound on the regret almost equivalent to (though slightly better than) that achieved by projection on the subset Δ_d^α of the simplex.

Theorem 2. With the above update, for all $T \geq 1$, for all sequences ℓ_1, \dots, ℓ_T of loss vectors $\ell_t \in [0, 1]^d$, and for all $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathbb{R}_+^d$,

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{u}_t\|_1 \widehat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \mathbf{u}_t^\top \ell_t &\leq \frac{\|\mathbf{u}_1\|_1 \ln d}{\eta} + \frac{\eta}{8} \sum_{t=1}^T \|\mathbf{u}_t\|_1 \\ &\quad + \frac{m(\mathbf{u}_1^T)}{\eta} \ln \frac{d}{\alpha} + \frac{\sum_{t=2}^T \|\mathbf{u}_t\|_1 - m(\mathbf{u}_1^T)}{\eta} \ln \frac{1}{1-\alpha}. \end{aligned}$$

Note that if we only consider vectors of the form $\mathbf{u}_t = \mathbf{q}_t = (0, \dots, 0, 1, 0, \dots, 0)$ then $m(\mathbf{q}_1^T)$ corresponds to the number of times $\mathbf{q}_{t+1} \neq \mathbf{q}_t$ in the sequence \mathbf{q}_1^T . We thus recover [4, Theorem 1] and [6, Lemma 6] from the much more general Theorem 2.

The fixed-share forecaster does not need to “know” anything in advance about the sequence of the norms $\|\mathbf{u}_t\|_1$ for the bound above to be valid. Of course, in order to minimize the obtained upper bound, the tuning parameters α, η need to be optimized and their values will depend on the maximal values of $m(\mathbf{u}_1^T)$ and $\sum_{t=1}^T \|\mathbf{u}_t\|_1$ for the sequences one wishes to compete against. This is illustrated in the following corollary, whose proof is omitted. Therein, $h(x) = -x \ln x - (1-x) \ln(1-x)$ denotes the binary entropy function for $x \in [0, 1]$. We recall³ that $h(x) \leq x \ln(e/x)$ for $x \in [0, 1]$.

Corollary 1. Suppose Algorithm 1 is run with the update (5). Let $m_0 > 0$ and $U_0 > 0$. For all $T \geq 1$, for all sequences ℓ_1, \dots, ℓ_T of loss vectors $\ell_t \in [0, 1]^d$, and for all sequences $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathbb{R}_+^d$ with $\|\mathbf{u}_1\|_1 + m(\mathbf{u}_1^T) \leq m_0$ and $\sum_{t=1}^T \|\mathbf{u}_t\|_1 \leq U_0$,

$$\sum_{t=1}^T \|\mathbf{u}_t\|_1 \widehat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \mathbf{u}_t^\top \ell_t \leq \sqrt{\frac{U_0}{2} \left(m_0 \ln d + U_0 h\left(\frac{m_0}{U_0}\right) \right)} \leq \sqrt{\frac{U_0 m_0}{2} \left(\ln d + \ln\left(\frac{e U_0}{m_0}\right) \right)}$$

whenever η and α are optimally chosen in terms of m_0 and U_0 .

Proof of Theorem 2. Applying Lemma 1 with $\mathbf{q}_t = \mathbf{u}_t / \|\mathbf{u}_t\|_1$, and multiplying by $\|\mathbf{u}_t\|_1$, we get for all $t \geq 1$ and $\mathbf{u}_t \in \mathbb{R}_+^d$

$$\|\mathbf{u}_t\|_1 \widehat{\mathbf{p}}_t^\top \ell_t - \mathbf{u}_t^\top \ell_t \leq \frac{1}{\eta} \sum_{i=1}^d u_{i,t} \ln \frac{v_{i,t+1}}{\widehat{p}_{i,t}} + \frac{\eta}{8} \|\mathbf{u}_t\|_1. \quad (6)$$

³As can be seen by noting that $\ln(1/(1-x)) < x/(1-x)$

We now examine

$$\sum_{i=1}^d u_{i,t} \ln \frac{v_{i,t+1}}{\widehat{p}_{i,t}} = \sum_{i=1}^d \left(u_{i,t} \ln \frac{1}{\widehat{p}_{i,t}} - u_{i,t-1} \ln \frac{1}{v_{i,t}} \right) + \sum_{i=1}^d \left(u_{i,t-1} \ln \frac{1}{v_{i,t}} - u_{i,t} \ln \frac{1}{v_{i,t+1}} \right). \quad (7)$$

For the first term on the right-hand side, we have

$$\begin{aligned} \sum_{i=1}^d \left(u_{i,t} \ln \frac{1}{\widehat{p}_{i,t}} - u_{i,t-1} \ln \frac{1}{v_{i,t}} \right) &= \sum_{i: u_{i,t} \geq u_{i,t-1}} \left((u_{i,t} - u_{i,t-1}) \ln \frac{1}{\widehat{p}_{i,t}} + u_{i,t-1} \ln \frac{v_{i,t}}{\widehat{p}_{i,t}} \right) \\ &\quad + \underbrace{\sum_{i: u_{i,t} < u_{i,t-1}} \left((u_{i,t} - u_{i,t-1}) \ln \frac{1}{v_{i,t}} + u_{i,t} \ln \frac{v_{i,t}}{\widehat{p}_{i,t}} \right)}_{\leq 0}. \end{aligned} \quad (8)$$

In view of the update (5), we have $1/\widehat{p}_{i,t} \leq d/\alpha$ and $v_{i,t}/\widehat{p}_{i,t} \leq 1/(1-\alpha)$. Substituting in (8), we get

$$\begin{aligned} &\sum_{i=1}^d \left(u_{i,t} \ln \frac{1}{\widehat{p}_{i,t}} - u_{i,t-1} \ln \frac{1}{v_{i,t}} \right) \\ &\leq \sum_{i: u_{i,t} \geq u_{i,t-1}} (u_{i,t} - u_{i,t-1}) \ln \frac{d}{\alpha} + \left(\sum_{i: u_{i,t} \geq u_{i,t-1}} u_{i,t-1} + \sum_{i: u_{i,t} < u_{i,t-1}} u_{i,t} \right) \ln \frac{1}{1-\alpha} \\ &= D_{\text{TV}}(\mathbf{u}_t, \mathbf{u}_{t-1}) \ln \frac{d}{\alpha} + \underbrace{\left(\sum_{i=1}^d u_{i,t} - \sum_{i: u_{i,t} \geq u_{i,t-1}} (u_{i,t} - u_{i,t-1}) \right)}_{= \|\mathbf{u}_t\|_1 - D_{\text{TV}}(\mathbf{u}_t, \mathbf{u}_{t-1})} \ln \frac{1}{1-\alpha}. \end{aligned}$$

The sum of the second term in (7) telescopes. Substituting the obtained bounds in the first sum of the right-hand side in (7), and summing over $t = 2, \dots, T$, leads to

$$\begin{aligned} \sum_{t=2}^T \sum_{i=1}^d u_{i,t} \ln \frac{v_{i,t+1}}{\widehat{p}_{i,t}} &\leq m(\mathbf{u}_1^T) \ln \frac{d}{\alpha} + \left(\sum_{t=2}^T \|\mathbf{u}_t\|_1 - m(\mathbf{u}_1^T) \right) \ln \frac{1}{1-\alpha} \\ &\quad + \sum_{i=1}^d u_{i,1} \ln \frac{1}{v_{i,2}} - \underbrace{u_{i,T} \ln \frac{1}{v_{i,T+1}}}_{\leq 0}. \end{aligned}$$

We hence get from (6), which we use in particular for $t = 1$,

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{u}_t\|_1 \widehat{\mathbf{p}}_t^\top \boldsymbol{\ell}_t - \mathbf{u}_t^\top \boldsymbol{\ell}_t &\leq \frac{1}{\eta} \sum_{i=1}^d u_{i,1} \ln \frac{1}{\widehat{p}_{i,1}} + \frac{\eta}{8} \sum_{t=1}^T \|\mathbf{u}_t\|_1 \\ &\quad + \frac{m(\mathbf{u}_1^T)}{\eta} \ln \frac{d}{\alpha} + \frac{\sum_{t=2}^T \|\mathbf{u}_t\|_1 m(\mathbf{u}_1^T)}{\eta} \ln \frac{1}{1-\alpha}. \quad \square \end{aligned}$$

6 Applications

We now show how our regret bounds can be specialized to obtain bounds on adaptive and discounted regret, and on regret with time-selection functions. We show regret bounds only for the specific instance of the generalized share algorithm using update (5); but the discussion below also holds up to minor modifications for the forecaster studied in Theorem 1.

Adaptive regret was introduced by [8] and can be viewed as a variant of discounted regret where the monotonicity assumption is dropped. For $\tau_0 \in \{1, \dots, T\}$, the τ_0 -*adaptive regret* of a forecaster is defined by

$$\mathcal{R}_T^{\tau_0\text{-adapt}} = \max_{\substack{[r, s] \subset [1, T] \\ s+1-r \leq \tau_0}} \left\{ \sum_{t=r}^s \widehat{\mathbf{p}}_t^\top \boldsymbol{\ell}_t - \min_{\mathbf{q} \in \Delta_d} \sum_{t=r}^s \mathbf{q}^\top \boldsymbol{\ell}_t \right\}. \quad (9)$$

The fact that this is a special case of (2) clearly emerges from the proof of Corollary 2 below here.

Adaptive regret is an alternative way to measure the performance of a forecaster against a changing environment. It is a straightforward observation that adaptive regret bounds also lead to shifting regret bounds (in terms of hard shifts). In this paper we note that these two notions of regret share an even tighter connection, as they can be both viewed as instances of the same *alma mater* notion of regret, i.e., the generalized shifting regret introduced in Section 3. The work [8] essentially considered the case of online convex optimization with exp-concave loss function; in case of general convex functions, they also mentioned that the greedy projection forecaster of [2] enjoys adaptive regret guarantees. This is obtained in much the same way as we obtain an adaptive regret bound for the fixed-share forecaster in the next result.

Corollary 2. *Suppose that Algorithm 1 is run with the shared update (5). Then for all $T \geq 1$, for all sequences ℓ_1, \dots, ℓ_T of loss vectors $\ell_t \in [0, 1]^d$, and for all $\tau_0 \in \{1, \dots, T\}$,*

$$\mathcal{R}_T^{\tau_0\text{-adapt}} \leq \sqrt{\frac{\tau_0}{2} \left(\tau_0 h\left(\frac{1}{\tau_0}\right) + \ln d \right)} \leq \sqrt{\frac{\tau_0}{2} \ln(ed\tau_0)}$$

whenever η and α are chosen optimally (depending on τ_0 and T).

As mentioned in [8], standard lower bounds on the regret show that the obtained bound is optimal up to the logarithmic factors.

Proof. For $1 \leq r \leq s \leq T$ and $\mathbf{q} \in \Delta_d$, the regret in the right-hand side of (9) equals the regret considered in Theorem 2 against the sequence \mathbf{u}_1^T defined as $\mathbf{u}_t = \mathbf{q}$ for $t = r, \dots, s$ and $\mathbf{0} = (0, \dots, 0)$ for the remaining t . When $r \geq 2$, this sequence is such that $D_{\text{TV}}(\mathbf{u}_r, \mathbf{u}_{r-1}) = D_{\text{TV}}(\mathbf{q}, \mathbf{0}) = 1$ and $D_{\text{TV}}(\mathbf{u}_{s+1}, \mathbf{u}_s) = D_{\text{TV}}(\mathbf{0}, \mathbf{q}) = 0$ so that $m(\mathbf{u}_1^T) = 1$, while $\|\mathbf{u}_1\|_1 = 0$. When $r = 1$, we have $\|\mathbf{u}_1\|_1 = 1$ and $m(\mathbf{u}_1^T) = 0$. In all cases, $m(\mathbf{u}_1^T) + \|\mathbf{u}_1\|_1 = 1$, that is, $m_0 = 1$. Specializing the bound of Theorem 2 with the additional choice $U_0 = \tau_0$ gives the result. \square

Discounted regret was introduced in [3, Section 2.11] and is defined by

$$\max_{\mathbf{q} \in \Delta_d} \sum_{t=1}^T \beta_{t,T} (\hat{\mathbf{p}}_t^\top \ell_t - \mathbf{q}^\top \ell_t). \quad (10)$$

The discount factors $\beta_{t,T}$ measure the relative importance of more recent losses to older losses. For instance, for a given horizon T , the discounts $\beta_{t,T}$ may be larger as t is closer to T . On the contrary, in a game-theoretic setting, the earlier losses may matter more than the more recent ones (because of interest rates), in which case $\beta_{t,T}$ would be smaller as t gets closer to T . We mostly consider below monotonic sequences of discounts (both non-decreasing and non-increasing). Up to a normalization, we assume that all discounts $\beta_{t,T}$ are in $[0, 1]$. As shown in [3], a minimal requirement to get non-trivial bounds is that the sum of the discounts satisfies $U_T = \sum_{t \leq T} \beta_{t,T} \rightarrow \infty$ as $T \rightarrow \infty$.

A natural objective is to show that the quantity in (10) is $o(U_T)$, for instance, by bounding it by something of the order of $\sqrt{U_T}$. We claim that Corollary 1 does so, at least whenever the sequences $(\beta_{t,T})$ are monotonic for all T . To support this claim, we only need to show that $m_0 = 1$ is a suitable value to deal with (10). Indeed, for all $T \geq 1$ and for all $\mathbf{q} \in \Delta_d$, the measure of regularity involved in the corollary satisfies

$$\|\beta_{1,T} \mathbf{q}\|_1 + m((\beta_{t,T} \mathbf{q})_{t \leq T}) = \beta_{1,T} + \sum_{t=2}^T (\beta_{t,T} - \beta_{t-1,T})_+ = \max\{\beta_{1,T}, \beta_{T,T}\} \leq 1,$$

where the second equality follows from the monotonicity assumption on the discounts.

The values of the discounts for all t and T are usually known in advance. However, the horizon T is not. Hence, a calibration issue may arise. The online tuning of the parameters α and η shown in Section 7.3 entails a forecaster that can get discounted regret bounds of the order $\sqrt{U_T}$ for all T . The fundamental reason for this is that the discounts only come in the definition of the fixed-share forecaster via their sums. In contrast, the forecaster discussed in [3, Section 2.11] weighs each instance t directly with $\beta_{t,T}$ (i.e., in the very definition of the forecaster) and enjoys therefore no regret guarantees for horizons other than T (neither before T nor after T). Therein, the knowledge

of the horizon T is so crucial that it cannot be dealt with easily, not even with online calibration of the parameters or with a doubling trick. We insist that for the fixed-share forecaster, much flexibility is gained as some of the discounts $\beta_{t,T}$ can change in a drastic manner for a round T to values $\beta_{t,T+1}$ for the next round. However we must admit that the bound of [3, Section 2.11] is smaller than the one obtained above, as it is of the order of $\sqrt{\sum_{t \leq T} \beta_{t,T}^2}$, in contrast to our $\sqrt{\sum_{t \leq T} \beta_{t,T}}$ bound. Again, this improvement was made possible because of the knowledge of the time horizon.

As for the comparison to the setting of discounted losses of [9], we note that the latter can be cast as a special case of our setting (since the discounting weights take the special form $\beta_{t,T} = \gamma_t \dots \gamma_{T-1}$ therein, for some sequence γ_s of positive numbers). In particular, the fixed-share forecaster can satisfy the bound stated in [9, Theorem 2], for instance, by using the online tuning techniques of Section 7.3. A final reference to mention is the setting of time-selection functions of [10, Section 6], which basically corresponds to knowing in advance the weights $\|\mathbf{u}_t\|_1$ of the comparison sequence $\mathbf{u}_1, \dots, \mathbf{u}_T$ the forecaster will be evaluated against. We thus generalize their results as well.

7 Refinements and extensions

We now show that techniques for refining the standard online analysis can be easily applied to our framework. We focus on the following: improvement for small losses, sparse target sequences, and dynamic tuning of parameters. Not all of them were within reach of previous analyses.

7.1 Improvement for small losses

The regret bounds of the fixed-share forecaster can be significantly improved when the cumulative loss of the best sequence of experts is small. The next result improves on Corollary 1 whenever $L_0 \ll U_0$. For concreteness, we focus on the fixed-share update (5).

Corollary 3. *Suppose Algorithm 1 is run with the update (5). Let $m_0 > 0$, $U_0 > 0$, and $L_0 > 0$. For all $T \geq 1$, for all sequences ℓ_1, \dots, ℓ_T of loss vectors $\ell_t \in [0, 1]^d$, and for all sequences $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathbb{R}_+^d$ with $\|\mathbf{u}_1\|_1 + m(\mathbf{u}_1^T) \leq m_0$, $\sum_{t=1}^T \|\mathbf{u}_t\|_1 \leq U_0$, and $\sum_{t=1}^T \mathbf{u}_t^\top \ell_t \leq L_0$,*

$$\sum_{t=1}^T \|\mathbf{u}_t\|_1 \widehat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \mathbf{u}_t^\top \ell_t \leq \sqrt{L_0 m_0 \left(\ln d + \ln \left(\frac{e U_0}{m_0} \right) \right)} + \ln d + \ln \left(\frac{e U_0}{m_0} \right)$$

whenever η and α are optimally chosen in terms of m_0 , U_0 , and L_0 .

Here again, the parameters α and η may be tuned online using the techniques shown in Section 7.3. The above refinement is obtained by mimicking the analysis of Hedge forecasters for small losses (see, e.g., [3, Section 2.4]). In particular, one should substitute Lemma 1 with the following lemma in the analysis carried out in Section 5; its proof follows from the mere replacement of Hoeffding's inequality by [3, Lemma A.3], which states that for all $\eta \in \mathbb{R}$ and for all random variable X taking values in $[0, 1]$, one has $\ln \mathbb{E}[e^{-\eta X}] \leq (e^{-\eta} - 1)\mathbb{E}X$.

Lemma 2. *Algorithm 1 satisfies $\frac{1 - e^{-\eta}}{\eta} \widehat{\mathbf{p}}_t^\top \ell_t - \mathbf{q}_t^\top \ell_t \leq \frac{1}{\eta} \sum_{i=1}^d q_{i,t} \ln \left(\frac{v_{i,t}}{\widehat{p}_{i,t+1}} \right)$ for all $\mathbf{q}_t \in \Delta_d$.*

7.2 Sparse target sequences

The work [6] introduced forecasters that are able to efficiently compete with the best sequence of experts among all those sequences that only switch a bounded number of times and also take a small number of different values. Such “sparse” sequences of experts appear naturally in many applications. In this section we show that their algorithms in fact work very well in comparison with a much larger class of sequences $\mathbf{u}_1, \dots, \mathbf{u}_T$ that are “regular”—that is, $m(\mathbf{u}_1^T)$, defined in (3) is small—and “sparse” in the sense that the quantity $n(\mathbf{u}_1^T) = \sum_{i=1}^d \max_{t=1, \dots, T} u_{i,t}$ is small. Note that when $\mathbf{q}_t \in \Delta_d$ for all t , then two interesting upper bounds can be provided. First, denoting the union of the supports of these convex combinations by $S \subseteq \{1, \dots, d\}$, we have $n(\mathbf{q}_1^T) \leq |S|$, the cardinality of S . Also, $n(\mathbf{q}_1^T) \leq |\{\mathbf{q}_t, t = 1, \dots, T\}|$, the cardinality of the pool of convex combinations. Thus, $n(\mathbf{u}_1^T)$ generalizes the notion of sparsity of [6].

Here we consider a family of shared updates of the form

$$\hat{p}_{j,t} = (1 - \alpha)v_{j,t} + \alpha \frac{w_{j,t}}{Z_t}, \quad 0 \leq \alpha \leq 1, \quad (11)$$

where the $w_{j,t}$ are nonnegative weights that may depend on past and current pre-weights and $Z_t = \sum_{i=1}^d w_{i,t}$ is a normalization constant. Shared updates of this form were proposed by [6, Sections 3 and 5.2]. Apart from generalizing the regret bounds of [6], we believe that the analysis given below is significantly simpler and more transparent. We are also able to slightly improve their original bounds.

We focus on choices of the weights $w_{j,t}$ that satisfy the following conditions: there exists a constant $C \geq 1$ such that for all $j = 1, \dots, d$ and $t = 1, \dots, T$,

$$v_{j,t} \leq w_{j,t} \leq 1 \quad \text{and} \quad C w_{j,t+1} \geq w_{j,t}. \quad (12)$$

The next result improves on Theorem 2 when $T \ll d$ and $n(\mathbf{u}_1^T) \ll m(\mathbf{u}_1^T)$, that is, when the dimension (or number of experts) d is large but the sequence \mathbf{u}_1^T is sparse. Its proof can be found in the supplementary material; it is a variation on the proof of Theorem 2.

Theorem 3. *Suppose Algorithm 1 is run with the shared update (11) with weights satisfying the conditions (12). Then for all $T \geq 1$, for all sequences ℓ_1, \dots, ℓ_T of loss vectors $\ell_t \in [0, 1]^d$, and for all sequences $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathbb{R}_+^d$,*

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{u}_t\|_1 \hat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \mathbf{u}_t^\top \ell_t &\leq \frac{n(\mathbf{u}_1^T) \ln d}{\eta} + \frac{n(\mathbf{u}_1^T) T \ln C}{\eta} + \frac{\eta}{8} \sum_{t=1}^T \|\mathbf{u}_t\|_1 \\ &\quad + \frac{m(\mathbf{u}_1^T)}{\eta} \ln \frac{\max_{t \leq T} Z_t}{\alpha} + \frac{\sum_{t=2}^T \|\mathbf{u}_t\|_1 - m(\mathbf{u}_1^T)}{\eta} \ln \frac{1}{1 - \alpha}. \end{aligned}$$

Corollaries 8 and 9 of [6] can now be generalized (and even improved); we do so—in the supplementary material—by showing two specific instances of the generic update (11) that satisfy (12).

7.3 Online tuning of the parameters

The forecasters studied above need their parameters η and α to be tuned according to various quantities, including the time horizon T . We show here how the trick of [11] of having these parameters vary over time can be extended to our setting. For the sake of concreteness we focus on the fixed-share update, i.e., Algorithm 1 run with the update (5). We respectively replace steps 3 and 4 of its description by the loss and shared updates

$$v_{j,t+1} = \hat{p}_{j,t}^{\frac{\eta_t}{\eta_{t-1}}} e^{-\eta_t \ell_{j,t}} \bigg/ \sum_{i=1}^d \hat{p}_{i,t}^{\frac{\eta_t}{\eta_{t-1}}} e^{-\eta_t \ell_{i,t}} \quad \text{and} \quad p_{j,t+1} = \frac{\alpha_t}{d} + (1 - \alpha_t) v_{j,t+1}, \quad (13)$$

for all $t \geq 1$ and all $j \in \{1, \dots, d\}$, where (η_τ) and (α_τ) are two sequences of positive numbers, indexed by $\tau \geq 1$. We also conventionally define $\eta_0 = \eta_1$. Theorem 2 is then adapted in the following way (when $\eta_t \equiv \eta$ and $\alpha_t \equiv \alpha$, Theorem 2 is exactly recovered).

Theorem 4. *The forecaster based on the updates (13) is such that whenever $\eta_t \leq \eta_{t-1}$ and $\alpha_t \leq \alpha_{t-1}$ for all $t \geq 1$, the following performance bound is achieved. For all $T \geq 1$, for all sequences ℓ_1, \dots, ℓ_T of loss vectors $\ell_t \in [0, 1]^d$, and for all $\mathbf{u}_1, \dots, \mathbf{u}_T \in \mathbb{R}_+^d$,*

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{u}_t\|_1 \hat{\mathbf{p}}_t^\top \ell_t - \sum_{t=1}^T \mathbf{u}_t^\top \ell_t &\leq \left(\frac{\|\mathbf{u}_t\|_1}{\eta_1} + \sum_{t=2}^T \|\mathbf{u}_t\|_1 \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \right) \ln d \\ &\quad + \frac{m(\mathbf{u}_1^T)}{\eta_T} \ln \frac{d(1 - \alpha_T)}{\alpha_T} + \sum_{t=2}^T \frac{\|\mathbf{u}_t\|_1}{\eta_{t-1}} \ln \frac{1}{1 - \alpha_t} + \sum_{t=1}^T \frac{\eta_{t-1}}{8} \|\mathbf{u}_t\|_1. \end{aligned}$$

Due to space constraints, we provide an illustration of this bound only in the supplementary material.

Acknowledgments

The authors acknowledge support from the French National Research Agency (ANR) under grant EXPLO/RA (“Exploration–exploitation for efficient resource allocation”) and by the PASCAL2 Network of Excellence under EC grant no. 506778.

References

- [1] M. Herbster and M. Warmuth. Tracking the best linear predictor. *Journal of Machine Learning Research*, 1:281–309, 2001.
- [2] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning, ICML 2003*, 2003.
- [3] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [4] M. Herbster and M. Warmuth. Tracking the best expert. *Machine Learning*, 32:151–178, 1998.
- [5] V. Vovk. Derandomizing stochastic prediction strategies. *Machine Learning*, 35(3):247–282, Jun. 1999.
- [6] O. Bousquet and M.K. Warmuth. Tracking a small set of experts by mixing past posteriors. *Journal of Machine Learning Research*, 3:363–396, 2002.
- [7] A. György, T. Linder, and G. Lugosi. Tracking the best of many experts. In *Proceedings of the 18th Annual Conference on Learning Theory (COLT)*, pages 204–216, Bertinoro, Italy, Jun. 2005. Springer.
- [8] E. Hazan and C. Seshadhri. Efficient learning algorithms for changing environments. *Proceedings of the 26th International Conference of Machine Learning (ICML)*, 2009.
- [9] A. Chernov and F. Zhdanov. Prediction with expert advice under discounted loss. In *Proceedings of the 21st International Conference on Algorithmic Learning Theory, ALT 2010*, pages 255–269. Springer, 2008.
- [10] A. Blum and Y. Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, 2007.
- [11] P. Auer, N. Cesa-Bianchi, and C. Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 64:48–75, 2002.