Lower bounds on individual sequence regret

Eyal Gofer¹ · Yishay Mansour^{2,3}

Received: 27 December 2014 / Accepted: 25 August 2015 / Published online: 25 September 2015
© The Author(s) 2015

Abstract In this work we lower bound the individual sequence anytime regret of a large family of online algorithms. This bound depends on the quadratic variation of the sequence, Q_T , and the learning rate. Nevertheless, we show that any learning rate that guarantees a regret upper bound of $O(\sqrt{Q_T})$ necessarily implies an $\Omega(\sqrt{Q_T})$ anytime regret on *any* sequence with quadratic variation Q_T . The algorithms we consider are online linear optimization forecasters whose weight vector at time $t + 1$ is the gradient of a concave potential function of cumulative losses at time t . We show that these algorithms include all linear Regularized Follow the Leader algorithms. We prove our result for the case of potentials with negative definite Hessians, and potentials for the best expert setting satisfying some natural regularity conditions. In the best expert setting, we give our result in terms of the translation-invariant *relative* quadratic variation. We apply our lower bounds to Randomized Weighted Majority and to linear cost Online Gradient Descent. We show that our analysis can be generalized to accommodate diverse measures of variation beside quadratic variation. We apply this generalized analysis to Online Gradient Descent with a regret upper bound that depends on the variance of losses.

Keywords Regret minimization · Online learning · Online linear optimization · Regret lower bounds · Regularized Follow the Leader

Editor: Gabor Lugosi.

✉ Eyal Gofer
eyalgofe@cs.huji.ac.il

Yishay Mansour
mansour.yishay@gmail.com

¹ The Rachel and Selim Benin School of Computer Science and Engineering, The Hebrew University, Givat Ram, Jerusalem 91904, Israel

² Microsoft Research, Herzliya, Israel

³ Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel

1 Introduction

For any sequence of losses, it is trivial to tailor an algorithm that has no regret on that particular sequence. The challenge and the great success of regret minimization algorithms lie in achieving low regret for *every* sequence. This bound may still depend on a measure of sequence smoothness, such as the quadratic variation or variance. The optimality of such regret upper bounds may be demonstrated by proving the existence of some “difficult” loss sequences. Their existence may be implied, for example, by stochastically generating sequences and proving a lower bound on the expected regret of any algorithm (see, e.g., [Cesa-Bianchi and Lugosi 2006](#)). This type of argument leaves open the possibility that the difficult sequences are, in some way, atypical or irrelevant to actual user needs. In this work we address this question by proving lower bounds on the regret of *any* individual sequence, in terms of its quadratic variation.

We first consider the related task of characterizing algorithms that have individual sequence *non-negative* regret. We focus our attention on linear forecasters that determine their next weight vector as a function of current cumulative losses. More specifically, if $\mathbf{L}_t \in \mathbb{R}^N$ is the cumulative loss vector at time t , and $\mathbf{x}_{t+1} \in \mathcal{K}$ is the next weight vector, then $\mathbf{x}_{t+1} = g(\mathbf{L}_t)$ for some continuous g . The algorithm then incurs a loss of $\mathbf{x}_{t+1} \cdot \mathbf{l}_{t+1}$, where $\mathbf{l}_t$ is the loss vector at time t . We show that such algorithms have individual sequence non-negative regret if and only if g is the gradient of a concave potential function. We then show that this characteristic is shared by all linear cost Regularized Follow the Leader (RFTL) regret minimization algorithms, which include Randomized Weighted Majority (RWM) and linear cost Online Gradient Descent (OGD).

As our main result, we prove a trade-off between the upper bound on an algorithm’s regret and a lower bound on its anytime regret, namely, its maximal regret for any prefix of the loss sequence. In particular, if the algorithm has a regret upper bound of $O(\sqrt{Q})$ for any sequence with quadratic variation Q , then it must have an $\Omega(\sqrt{Q})$ anytime regret on any sequence with quadratic variation $\Theta(Q)$.

We prove our result for two separate classes of continuously twice-differentiable potentials. One class has negative definite Hessians in a neighborhood of $\mathbf{L} = \mathbf{0}$, and includes OGD. The other comprises potentials for the best expert setting whose Hessians in a neighborhood of $\mathbf{L} = \mathbf{0}$ have positive off-diagonal entries; in other words, such potentials increase the weights of experts as their performance relatively improves, which is a natural property of regret minimization algorithms. For the first class, we use the usual Euclidean quadratic variation, or $\sum_{t=1}^T \|\mathbf{l}_t\|_2^2$. For the best expert setting, however, we use the more appropriate *relative quadratic variation*, $\sum_{t=1}^T (\max_i \{l_{i,t}\} - \min_i \{l_{i,t}\})^2$.

Our proof is comprised of several steps. We add a learning rate η to any potential Φ by defining a new potential $\Phi_\eta(\mathbf{L}) = (1/\eta)\Phi(\eta\mathbf{L})$. We give an exact expression for the regret using the Taylor expansion of the potential, and use it to prove an $\Omega(\min\{1/\eta, \eta Q\})$ lower bound on the anytime regret of any sequence with quadratic variation Q . In addition, we construct two specific loss sequences with variation Q , one with an $\Omega(1/\eta)$ regret, and the other with $\Omega(\eta Q)$ regret. Thus, we must have $\eta = \Theta(1/\sqrt{Q})$ to ensure a regret of $O(\sqrt{Q})$ for every sequence with variation Q , and our lower bound on the anytime regret becomes $\Omega(\sqrt{Q})$. We demonstrate our result on RWM, as an example of a best expert potential, and on linear cost OGD, as an example of a potential with a negative definite Hessian.

We show that our analysis may be generalized to accommodate diverse notions of variation instead of quadratic variation. Our lower bounds on the anytime regret and our trade-off result in fact apply to any notion of variation V that is upper bounded for any loss sequence by

a constant times the quadratic variation, namely $V \leq \beta Q$. This condition is satisfied by key notions of variation featured in the regret minimization literature, notably Hazan and Kale (2010), Chiang et al. (2012). We apply this generalized analysis to the OGD algorithm in the online linear optimization setting, with V taken as the variance-like measure $\text{Var}_T = \sum_{t=1}^T \|\mathbf{l}_t - \boldsymbol{\mu}_T\|_2^2$ (where $\boldsymbol{\mu}_T$ is the mean of the losses). We obtain in particular that a learning rate that guarantees a regret upper bound of $O(\sqrt{V})$ for any sequence with $\text{Var}_T \leq V$ necessarily implies an $\Omega(\sqrt{\text{Var}_T})$ anytime regret for any sequence with $\text{Var}_T = \Theta(V)$.

We remark that our lower bounds are not applicable to every online learning algorithm or scenario. For instance, in the best expert setting, algorithms that set their learning rate as a function of the cumulative loss of the best expert may circumvent our lower bounds. In Sects. 3 and 4 we give a precise definition of the relevant setting and in Sect. 7.1 we further discuss the scope of our results.

Finally, we consider extensions to our results and highlight two applications that demonstrate the usability of our work. We will briefly mention here the first of these applications, which was included in the conference version of this work, and pertains to the financial problem of option pricing. Specifically, by leveraging our lower bounds on the anytime regret of RWM it is possible to obtain a lower bound on the price of “at the money” European call options in an arbitrage-free market.¹ The price lower bound obtained for such options is $\exp(0.1\sqrt{Q}) - 1 \approx 0.1\sqrt{Q}$ for $Q < 0.5$, where Q is the assumed quadratic variation of the stock’s log price ratios. This bound has the same asymptotic behavior as the Black-Scholes price, which is approximately $\sqrt{Q}/\sqrt{2\pi} \approx 0.4\sqrt{Q}$, and improves on a previous result by DeMarzo et al. (2006), who gave a $Q/10$ lower bound for $Q < 1$. For an in-depth treatment of the financial aspect of this work, see Gofer and Mansour (2012), Gofer (2014b).

1.1 Related work

There are numerous results providing worst case upper bounds for regret minimization algorithms and showing their optimality (see Cesa-Bianchi and Lugosi 2006). In the best expert setting, RWM has been shown to have an optimal regret of $O(\sqrt{T \log N})$. In the online convex optimization paradigm (Zinkevich 2003), upper bounds of $O(\sqrt{T})$ have been shown, along with $\Omega(\sqrt{T})$ lower bounds for linear cost functions (Hazan 2006). In both cases, regret lower bounds are proved by invoking a stochastic adversary and lower bounding the expected regret. Upper bounds based on various notions of quadratic variation and variance are given in Cesa-Bianchi et al. (2007) for the expert setting, in Hazan and Kale (2010) for online linear optimization and the expert setting in particular, and in Chiang et al. (2012) for the more general setting of online convex optimization.

A trade-off result is given in Even-Dar et al. (2008), where it is shown that a best expert algorithm with $O(\sqrt{T})$ regret must have a worst case $\Omega(\sqrt{T})$ regret to any fixed average of experts. To the best of our knowledge, there are no results that lower bound the regret on loss sequences in terms of their individual quadratic variation.

1.2 Outline

In Sect. 2 we provide notation and definitions. Section 3 characterizes algorithms with non-negative individual sequence regret, proves they include linear cost RFTL, and provides basic

¹ A European call option is a financial instrument that pays its holder at time T the sum of $\max\{S_T - K, 0\}$, where S_t is the price of a given stock at time t , and K is a set price. The option is termed “at the money” if $K = S_0$. A market is arbitrage-free if no algorithm trading in financial assets can guarantee profit without any risk of losing money.

regret lower bounds. Section 4 presents our main result on the trade-off between upper bounds on regret and lower bounds on individual sequence anytime regret. In Sect. 5 we apply our bounds to linear cost OGD and to RWM. In Sect. 6 we go beyond quadratic variation and generalize our analysis to other popular notions of variation. In Sect. 7 we consider extensions and describe applications of our results.

A shorter version of this work was published in the proceedings of The Twenty-Third International Conference on Algorithmic Learning Theory (ALT 2012). The present version contains new additions that are given in Sects. 6 and 7. A key new contribution is the generalization of previous analysis to diverse notions of variation beyond quadratic variation.

2 Preliminaries

2.1 Regret minimization

In the *best expert* setting, there are N available experts, and at each time step $1 \leq t \leq T$, an online algorithm A selects a distribution $\mathbf{p}_t$ over the N experts. After the choice is made, an adversary selects a loss vector $\mathbf{l}_t = (l_{1,t}, \dots, l_{N,t}) \in \mathbb{R}^N$, and the algorithm experiences a loss of $l_{A,t} = \mathbf{p}_t \cdot \mathbf{l}_t$. We write $L_{i,t} = \sum_{\tau=1}^t l_{i,\tau}$ for the cumulative loss of expert i at time t , $\mathbf{L}_t = (L_{1,t}, \dots, L_{N,t})$, and $L_{A,t} = \sum_{\tau=1}^t l_{A,\tau}$ for the cumulative loss of A at time t . The *regret* of A at time T is $R_{A,T} = L_{A,T} - \min_j \{L_{j,T}\}$. The aim of a regret minimization algorithm is to achieve small regret regardless of the loss vectors chosen by the adversary. The *anytime regret* of A is the maximal regret over time, namely, $\max_t \{R_{A,t}\}$. We will sometimes use the notation $m(t) = \arg \min_i \{L_{i,t}\}$, where we take the smallest such index in case of a tie. Similarly, we write $M(t)$ for the expert with the maximal cumulative loss at time t . We also use the notation $\delta(\mathbf{v}) = \max_i \{v_i\} - \min_i \{v_i\}$ for any $\mathbf{v} \in \mathbb{R}^N$, so $\delta(\mathbf{L}_t) = L_{M(t),t} - L_{m(t),t}$.

The best expert setting is a special case of the more general setting of *online linear optimization*. In this setting, at time t the algorithm, or *linear forecaster*, chooses a weight vector $\mathbf{x}_t \in \mathbb{R}^N$, and incurs a loss of $\mathbf{x}_t \cdot \mathbf{l}_t$. In this paper we assume that the weight vectors are chosen from a compact and convex set $\mathcal{K}$. The regret of a linear forecaster A is then defined as $R_{A,T} = L_{A,T} - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}_T\}$. The best expert setting is simply the case where $\mathcal{K} = \Delta_N$, the probability simplex over N elements.

2.1.1 Quadratic variation

We define the *quadratic variation* of the loss sequence $\mathbf{l}_1, \dots, \mathbf{l}_T$ as $Q_T = \sum_{t=1}^T \|\mathbf{l}_t\|_2^2$. For the best expert setting, we will use the slightly different notion of *relative quadratic variation*, defined as $q_T = \sum_{t=1}^T \delta(\mathbf{l}_t)^2$. We write Q for a known lower bound on Q_T and q for a known lower bound on q_T .

2.2 Convex functions

We mention here some basic facts about convex and concave functions that we will require. For more on convex analysis, see Rockafellar (1970), Boyd and Vandenberghe (2004), and Nesterov (2004), among others.

We will discuss functions defined on $\mathbb{R}^N$. A function $f : C \rightarrow \mathbb{R}$ is *convex*, if C is a convex set and if for every $\lambda \in [0, 1]$ and $\mathbf{x}, \mathbf{y} \in C$, $f(\lambda\mathbf{x} + (1-\lambda)\mathbf{y}) \leq \lambda f(\mathbf{x}) + (1-\lambda)f(\mathbf{y})$. The function f is *concave* if $-f$ is convex. The function f is *strictly convex* if the inequality is strict for $\mathbf{x} \neq \mathbf{y}$ and $\lambda \in (0, 1)$. The function f is *strongly convex* with parameter $\alpha > 0$,

if for every $\mathbf{x}, \mathbf{y} \in C$ and $\lambda \in [0, 1]$, $f(\lambda\mathbf{x} + (1 - \lambda)\mathbf{y}) \leq \lambda f(\mathbf{x}) + (1 - \lambda)f(\mathbf{y}) - (\alpha/2)\lambda(1 - \lambda)\|\mathbf{x} - \mathbf{y}\|_2^2$. If f is differentiable on a convex set C , then f is convex iff for every $\mathbf{x}, \mathbf{y} \in C$, $\nabla f(\mathbf{y}) \cdot (\mathbf{y} - \mathbf{x}) \geq f(\mathbf{y}) - f(\mathbf{x}) \geq \nabla f(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x})$; f is strictly convex iff the above inequalities are strict for $\mathbf{x} \neq \mathbf{y}$. If f is twice differentiable, then it is convex iff its Hessian is positive semi-definite: for every $\mathbf{x} \in C$, $\nabla^2 f(\mathbf{x}) \succeq 0$. The *convex conjugate* of f (defined on $\text{dom} f$) is the function $f^*(\mathbf{y}) = \sup_{\mathbf{x} \in \text{dom} f} \{\mathbf{x} \cdot \mathbf{y} - f(\mathbf{x})\}$, which is convex, and its effective domain is $\text{dom} f^* = \{\mathbf{y} : f^*(\mathbf{y}) < \infty\}$.

2.3 Seminorms

A *seminorm* on $\mathbb{R}^N$ is a function $\|\cdot\| : \mathbb{R}^N \rightarrow \mathbb{R}$ with the following properties:

- Positive homogeneity: for every $a \in \mathbb{R}$, $\mathbf{x} \in \mathbb{R}^N$, $\|a\mathbf{x}\| = |a|\|\mathbf{x}\|$.
- Triangle inequality: for every $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^N$, $\|\mathbf{x} + \mathbf{x}'\| \leq \|\mathbf{x}\| + \|\mathbf{x}'\|$.

Clearly, every norm is a seminorm. A seminorm satisfies $\|\mathbf{x}\| \geq 0$ for every $\mathbf{x}$, and $\|\mathbf{0}\| = 0$. However, unlike a norm, $\|\mathbf{x}\| = 0$ does not imply $\mathbf{x} = \mathbf{0}$. We will not deal with the trivial all-zero seminorm. Thus, there always exists a vector with non-zero seminorm, and by homogeneity, there exists a vector with seminorm a for any $a \in \mathbb{R}^+$.

2.4 Miscellaneous notation

For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N$, we write $[\mathbf{x}, \mathbf{y}]$ for the line segment between $\mathbf{x}$ and $\mathbf{y}$, namely, $\{a\mathbf{x} + (1-a)\mathbf{y} : 0 \leq a \leq 1\}$. We use the notation $\text{conv}(A)$ for the convex hull of a set $A \subseteq \mathbb{R}^N$, that is, $\text{conv}(A) = \{\sum_{i=1}^k \lambda_i \mathbf{x}_i : \mathbf{x}_i \in A, \lambda_i \geq 0, i = 1, \dots, k, \sum_{i=1}^k \lambda_i = 1, k \in \mathbb{N}\}$.

3 Non-negative individual sequence regret

Our ultimate goal is to prove strictly positive individual sequence regret lower bounds for a variety of algorithms. In this section, we will characterize algorithms for which this goal is achievable, and prove some basic regret lower bounds. This will be done by considering the larger family of algorithms that have *non-negative* regret for any loss sequence. This family, it turns out, can be characterized exactly, and includes the important class of linear cost RFTL algorithms.

We focus on linear forecasters whose vector at time t is determined as $\mathbf{x}_t = g(\mathbf{L}_{t-1})$, for $1 \leq t \leq T$, where $g : \mathbb{R}^N \rightarrow \mathcal{K} \subset \mathbb{R}^N$ is continuous and $\mathcal{K}$ is compact and convex. For such algorithms we can write $L_{A,T} = \sum_{t=1}^T g(\mathbf{L}_{t-1}) \cdot (\mathbf{L}_t - \mathbf{L}_{t-1})$.

A *non-negative-regret algorithm* satisfies that $L_{A,T} \geq \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\}$ for every $\mathbf{L}_0, \dots, \mathbf{L}_T \in \mathbb{R}^N$. We point out that we allow both positive and negative losses. Synonymously, we will also say that g has non-negative regret. Note that if $\mathbf{L}_0 = \mathbf{L}_T$ (a closed cumulative loss path), then for any $\mathbf{u} \in \mathcal{K}$, it holds that $\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0) = 0$, and non-negative regret implies that $\sum_{t=1}^T g(\mathbf{L}_{t-1}) \cdot (\mathbf{L}_t - \mathbf{L}_{t-1}) \geq 0$.

The following theorem gives an exact characterization of non-negative-regret forecasters as the gradients of concave potentials.²

Theorem 1 *A linear forecaster based on a continuous function g has individual sequence non-negative regret iff there exists a concave potential function $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}$ s.t. $g = \nabla \Phi$.*

² See the closely related Theorem 24.8 in Rockafellar (1970) regarding cyclically monotone mappings. The proof we give is different in that it involves the regret and relates the loss to the path integral.

Proof Let g have non-negative regret. For any path $\gamma : [0, 1] \rightarrow \mathbb{R}^N$ with finite length, we denote by $\int_\gamma \sum_{j=1}^N g_j dL_j$ the path integral of the vector field g along γ . For a closed path γ , we have that

$$\begin{aligned} \int_\gamma \sum_{j=1}^N g_j dL_j &= \lim_{n \rightarrow \infty} \sum_{i=1}^n g(\gamma((i-1)/n)) \cdot (\gamma(i/n) - \gamma((i-1)/n)) \\ &\geq \lim_{n \rightarrow \infty} \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\gamma(1) - \gamma(0))\} = 0, \end{aligned}$$

where the inequality is true because g has non-negative regret, and the last equality is true because $\gamma(1) = \gamma(0)$. For the reverse path γ_r , defined by $\gamma_r(t) = \gamma(1-t)$, we have $-\int_\gamma \sum_{j=1}^N g_j dL_j = \int_{\gamma_r} \sum_{j=1}^N g_j dL_j \geq 0$, where the inequality is true by the fact that γ_r is also a closed path. We thus have that $\int_\gamma \sum_{j=1}^N g_j dL_j = 0$ for any closed path, and therefore, g is a conservative vector field. As such, it has a potential function Φ s.t. $g = \nabla \Phi$.

Next, consider the case $\mathbf{L}_0 = \mathbf{L}, \mathbf{L}_1 = \mathbf{L}', \mathbf{L}_2 = \mathbf{L}$, for some arbitrary $\mathbf{L}, \mathbf{L}' \in \mathbb{R}^N$. Since g has non-negative regret, we have that $0 \leq g(\mathbf{L}) \cdot (\mathbf{L}' - \mathbf{L}) + g(\mathbf{L}') \cdot (\mathbf{L} - \mathbf{L}') = (g(\mathbf{L}) - g(\mathbf{L}')) \cdot (\mathbf{L}' - \mathbf{L})$. Therefore, $(\nabla \Phi(\mathbf{L}) - \nabla \Phi(\mathbf{L}')) \cdot (\mathbf{L} - \mathbf{L}') \leq 0$ for any $\mathbf{L}, \mathbf{L}' \in \mathbb{R}^N$, which means that Φ is concave (see, e.g., [Nesterov 2004](#), Theorem 2.1.3).

In the other direction, let $\mathbf{L}_0, \dots, \mathbf{L}_T \in \mathbb{R}^N$. We have that

$$\begin{aligned} \sum_{t=1}^T g(\mathbf{L}_{t-1}) \cdot (\mathbf{L}_t - \mathbf{L}_{t-1}) &= \sum_{t=1}^T \nabla \Phi(\mathbf{L}_{t-1}) \cdot (\mathbf{L}_t - \mathbf{L}_{t-1}) \geq \sum_{t=1}^T (\Phi(\mathbf{L}_t) - \Phi(\mathbf{L}_{t-1})) \\ &= \Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) \geq \nabla \Phi(\mathbf{L}_T) \cdot (\mathbf{L}_T - \mathbf{L}_0) \\ &\geq \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\}, \end{aligned}$$

where the first two inequalities are by the concavity of Φ . $\square$

As a by-product of the proof, we get the following:

Corollary 1 *If algorithm A uses a non-negative-regret function $g(\mathbf{L}) = \nabla \Phi(\mathbf{L})$, then it holds that $R_{A,T} \geq \Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\} \geq 0$.*

If Φ is continuously twice-differentiable, its second order Taylor expansion may be used to derive a similar lower bound, which now includes a non-negative quadratic regret term.

Theorem 2 *Let A be an algorithm using a non-negative-regret function $g = \nabla \Phi$, and let $\mathbf{L}_0, \dots, \mathbf{L}_T \in \mathbb{R}^N$. If Φ is continuously twice-differentiable on the set $\text{conv}(\{\mathbf{L}_0, \dots, \mathbf{L}_T\})$, then*

$$R_{A,T} = \Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\} - \frac{1}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\mathbf{z}_t) \mathbf{l}_t,$$

where $\mathbf{z}_t \in [\mathbf{L}_{t-1}, \mathbf{L}_t]$.

Proof For every $1 \leq t \leq T$, we have $\Phi(\mathbf{L}_t) - \Phi(\mathbf{L}_{t-1}) = \nabla \Phi(\mathbf{L}_{t-1}) \cdot \mathbf{l}_t + (1/2) \mathbf{l}_t^\top \nabla^2 \Phi(\mathbf{z}_t) \mathbf{l}_t$ for some $\mathbf{z}_t \in [\mathbf{L}_{t-1}, \mathbf{L}_t]$. Summing over $1 \leq t \leq T$, we get

$$\begin{aligned} \Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) &= \sum_{t=1}^T \nabla \Phi(\mathbf{L}_{t-1}) \cdot \mathbf{l}_t + \frac{1}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\mathbf{z}_t) \mathbf{l}_t \\ &= L_{A,T} + \frac{1}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\mathbf{z}_t) \mathbf{l}_t. \end{aligned}$$

Therefore,

$$\begin{aligned} R_{A,T} &= L_{A,T} - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\} \\ &= \Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\} - \frac{1}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\mathbf{z}_t) \mathbf{l}_t. \end{aligned}$$

□

Note that the regret is the sum of two non-negative terms. We will sometimes refer to the first one, $\Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\}$, as *the first order regret term*, and to the second, $-\frac{1}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\mathbf{z}_t) \mathbf{l}_t$, as *the second order regret term*. The second order term is non-negative by the concavity of Φ . These two terms play a key role in our bounds.

3.1 Relation to Regularized Follow the Leader

The class of concave potential algorithms contains the important class of linear cost RFTL algorithms. (For an in-depth discussion of the RFTL algorithm, see Hazan 2011.) Following Hazan (2011), let $\mathcal{K} \subset \mathbb{R}^N$ be compact, non-empty, and convex. $RFTL(\eta, \mathcal{R})$ is a linear forecaster with two parameters: the learning rate, $\eta > 0$, and a strongly convex and continuously twice-differentiable regularizing function, $\mathcal{R} : \mathcal{K} \rightarrow \mathbb{R}$. $RFTL(\eta, \mathcal{R})$ determines its weights according to the rule $\mathbf{x}_{t+1} = g(\mathbf{L}_t) = \arg \min_{\mathbf{x} \in \mathcal{K}} \{\mathbf{x} \cdot \mathbf{L}_t + \mathcal{R}(\mathbf{x})/\eta\}$.

The next theorem shows that linear RFTL is a concave potential algorithm, with a potential function that is directly related to the convex conjugate of the regularizing function. These are known properties (see, e.g., Lemma 15 in Shalev-Shwartz 2007), and a proof, which uses basic calculus, is given in Appendix 2 for completeness.

Theorem 3 *If $\mathcal{R} : \mathcal{K} \rightarrow \mathbb{R}$ is continuous and strongly convex and $\eta > 0$, then $\Phi(\mathbf{L}) = (-1/\eta)\mathcal{R}^*(-\eta\mathbf{L})$ is concave and continuously differentiable on $\mathbb{R}^N$, and for every $\mathbf{L} \in \mathbb{R}^N$, it holds that $\nabla \Phi(\mathbf{L}) = \arg \min_{\mathbf{x} \in \mathcal{K}} \{\mathbf{x} \cdot \mathbf{L} + \mathcal{R}(\mathbf{x})/\eta\}$ and $\Phi(\mathbf{L}) = \min_{\mathbf{x} \in \mathcal{K}} \{\mathbf{x} \cdot \mathbf{L} + \mathcal{R}(\mathbf{x})/\eta\}$.*

It is now possible to lower bound the regret of $RFTL(\eta, \mathcal{R})$ by applying the lower bounds of Corollary 1 and Theorem 2.

Theorem 4 *The regret of $RFTL(\eta, \mathcal{R})$ satisfies*

$$R_{RFTL(\eta, \mathcal{R}), T} \geq \frac{1}{\eta} (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1)) + \left(\mathbf{x}_{T+1} \cdot \mathbf{L}_T - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}_T\} \right) \geq 0,$$

and if $\mathcal{R}^*$ is continuously twice-differentiable on $\text{conv}(\{-\eta\mathbf{L}_0, \dots, -\eta\mathbf{L}_T\})$, then

$$\begin{aligned} R_{RFTL(\eta, \mathcal{R}), T} &\geq \frac{1}{\eta} (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1)) + \left(\mathbf{x}_{T+1} \cdot \mathbf{L}_T - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}_T\} \right) \\ &\quad + \frac{\eta}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \mathcal{R}^*(-\eta\mathbf{z}_t) \mathbf{l}_t \geq 0, \end{aligned}$$

where $\mathbf{z}_t \in [\mathbf{L}_{t-1}, \mathbf{L}_t]$.

Proof Let Φ be the potential function of $RFTL(\eta, \mathcal{R})$ according to Theorem 3. We have that $\Phi(\mathbf{L}_t) = \mathbf{x}_{t+1} \cdot \mathbf{L}_t + \mathcal{R}(\mathbf{x}_{t+1})/\eta$ and $\mathbf{x}_{t+1} = \nabla \Phi(\mathbf{L}_t)$ for every t . Therefore,

$$\begin{aligned}\Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) &= \mathbf{x}_{T+1} \cdot \mathbf{L}_T + \frac{1}{\eta} \mathcal{R}(\mathbf{x}_{T+1}) - \mathbf{x}_1 \cdot \mathbf{L}_0 - \frac{1}{\eta} \mathcal{R}(\mathbf{x}_1) \\ &= \frac{1}{\eta} (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1)) + \mathbf{x}_{T+1} \cdot \mathbf{L}_T,\end{aligned}$$

where we used the fact that $\mathbf{L}_0 = \mathbf{0}$. Therefore,

$$\begin{aligned}R_{RFTL(\eta, \mathcal{R}), T} &\geq \Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}_T\} \\ &= \frac{1}{\eta} (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1)) + \left(\mathbf{x}_{T+1} \cdot \mathbf{L}_T - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}_T\} \right) \geq 0,\end{aligned}$$

where the first inequality is by Corollary 1, and the second inequality is by the fact that $\mathbf{x}_1 = \arg \min_{\mathbf{u} \in \mathcal{K}} \{\mathcal{R}(\mathbf{u})\}$ and $\mathbf{x}_{T+1} \in \mathcal{K}$. This concludes the first part of the proof.

By Theorem 3, $\Phi(\mathbf{L}) = -(1/\eta)\mathcal{R}^*(-\eta\mathbf{L})$. Thus, Φ is continuously twice-differentiable on $\text{conv}(\{\mathbf{L}_0, \dots, \mathbf{L}_T\})$, and we have $\nabla^2 \Phi(\mathbf{L}) = -\eta \nabla^2 \mathcal{R}^*(-\eta\mathbf{L})$. By the first part and by Theorem 2, for some $\mathbf{z}_t \in [\mathbf{L}_{t-1}, \mathbf{L}_t]$,

$$\begin{aligned}R_{RFTL(\eta, \mathcal{R}), T} &= \Phi(\mathbf{L}_T) - \Phi(\mathbf{L}_0) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot (\mathbf{L}_T - \mathbf{L}_0)\} - \frac{1}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\mathbf{z}_t) \mathbf{l}_t \\ &= \frac{1}{\eta} (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1)) + \mathbf{x}_{T+1} \cdot \mathbf{L}_T - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}_T\} \\ &\quad + \frac{\eta}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \mathcal{R}^*(-\eta\mathbf{z}_t) \mathbf{l}_t,\end{aligned}$$

concluding the second part of the proof. $\square$

Note that the first order regret term is split into two new non-negative terms, namely, $(\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1))/\eta$ and $\mathbf{x}_{T+1} \cdot \mathbf{L}_T - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}_T\}$.

The first order regret lower bound given in Theorem 4 may be proved directly by extending the “follow the leader, be the leader” (FTL–BTL) lemma (Kalai and Vempala 2005, see also Hazan 2011). The extension and proof are given in Appendix 3.

4 Strictly positive individual sequence anytime regret

In this section we give lower bounds on the anytime regret for two classes of potentials. These are the class of potentials with negative definite Hessians, and a rich class of best expert potentials that includes RWM.

We start by describing our general argument, which is based on the lower bounds of Theorem 2 and Corollary 1. Note first that these bounds hold for any $1 \leq T' \leq T$. Now, let C be some convex set on which Φ is continuously twice-differentiable. If $\mathbf{L}_0, \dots, \mathbf{L}_T \in C$, then we may lower bound the regret by the second order term of Theorem 2, and further lower bound that term over any $\mathbf{L}_0, \dots, \mathbf{L}_T \in C$. Otherwise, there is some $1 \leq T' \leq T$ s.t. $\mathbf{L}_{T'} \notin C$, and we may lower bound the regret at time T' by the first order term, using Corollary 1. We further lower bound that term over any $\mathbf{L}_{T'} \notin C$. The minimum of those two lower bounds gives a lower bound on the anytime regret. By choosing C properly, we will be able to prove that these two lower bounds are strictly positive.

We now present our analysis in more detail. Observe that for every $\eta > 0$ and concave potential $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}$, $\Phi_\eta(\mathbf{L}) = (1/\eta)\Phi(\eta\mathbf{L})$ is also concave, with $\nabla \Phi_\eta(\mathbf{L}) = \nabla \Phi(\eta\mathbf{L})$.

Let $\|\cdot\|$ be a non-trivial seminorm on $\mathbb{R}^N$. In addition, let $a > 0$ be such that Φ is continuously twice-differentiable on the set $\{\|\mathbf{L}\| \leq a\}$, and let $\mathbf{L}_0 = \mathbf{0}$.

Suppose algorithm A uses Φ_η and encounters the loss path $\mathbf{L}_0, \dots, \mathbf{L}_T$. If there is some T' s.t. $\|\mathbf{L}_{T'}\| > a/\eta$, then applying Corollary 1 to Φ_η , we get

$$R_{A,T'} \geq \frac{1}{\eta} \left(\Phi(\eta \mathbf{L}_{T'}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \eta \mathbf{L}_{T'}\} \right).$$

Defining $\rho_1(a) = \inf_{\|\mathbf{L}\| \geq a} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\}$, we have that $R_{A,T'} \geq \rho_1(a)/\eta$.

We next assume that $\|\mathbf{L}_t\| \leq a/\eta$ for every t . It is easily verified that the set $\{\|\mathbf{L}\| \leq a/\eta\}$ is convex, and since it contains $\mathbf{L}_t$ for every t , it also contains $\text{conv}(\{\mathbf{L}_0, \dots, \mathbf{L}_T\})$. This means that $\Phi_\eta(\mathbf{L}) = (1/\eta)\Phi(\eta\mathbf{L})$ is continuously twice-differentiable on $\text{conv}(\{\mathbf{L}_0, \dots, \mathbf{L}_T\})$. Applying Theorem 2 to Φ_η and dropping the non-negative first order term, we have

$$R_{A,T} \geq -\frac{1}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi_\eta(\mathbf{z}_t) \mathbf{l}_t = -\frac{\eta}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\eta \mathbf{z}_t) \mathbf{l}_t,$$

where $\mathbf{z}_t \in [\mathbf{L}_{t-1}, \mathbf{L}_t]$. We now define $\rho_2(a) = \inf_{\|\mathbf{L}\| \leq a, \|\mathbf{l}\|=1} \{-\mathbf{l}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}\}$. If $\|\mathbf{l}_t\| \neq 0$, then

$$-\mathbf{l}_t^\top \nabla^2 \Phi(\eta \mathbf{z}_t) \mathbf{l}_t = -(\mathbf{l}_t / \|\mathbf{l}_t\|)^\top \nabla^2 \Phi(\eta \mathbf{z}_t) (\mathbf{l}_t / \|\mathbf{l}_t\|) \|\mathbf{l}_t\|^2 \geq \rho_2(a) \|\mathbf{l}_t\|^2,$$

where we used the fact that $\|\eta \mathbf{z}_t\| \leq a$, which holds since $\mathbf{z}_t \in \text{conv}(\{\mathbf{L}_0, \dots, \mathbf{L}_T\})$. Otherwise, $-\mathbf{l}_t^\top \nabla^2 \Phi(\eta \mathbf{z}_t) \mathbf{l}_t \geq 0 = \rho_2(a) \|\mathbf{l}_t\|^2$, so in any case,

$$R_{A,T} \geq -\frac{\eta}{2} \sum_{t=1}^T \mathbf{l}_t^\top \nabla^2 \Phi(\eta \mathbf{z}_t) \mathbf{l}_t \geq \frac{\eta}{2} \sum_{t=1}^T \rho_2(a) \|\mathbf{l}_t\|^2 = \frac{\eta}{2} \rho_2(a) \sum_{t=1}^T \|\mathbf{l}_t\|^2.$$

Thus, we have

Lemma 1 *If $\|\mathbf{L}_t\| \leq a/\eta$ for every t , then $R_{A,T} \geq \frac{\eta}{2} \rho_2(a) \sum_{t=1}^T \|\mathbf{l}_t\|^2$. Otherwise, for any t s.t. $\|\mathbf{L}_t\| > a/\eta$, $R_{A,t} \geq \rho_1(a)/\eta$. Therefore,*

$$\max_t \{R_{A,t}\} \geq \min \left\{ \frac{\rho_1(a)}{\eta}, \frac{\eta}{2} \rho_2(a) \sum_{t=1}^T \|\mathbf{l}_t\|^2 \right\}.$$

Note that $\rho_1(a)$ is non-decreasing, $\rho_2(a)$ is non-increasing, and $\rho_1(a), \rho_2(a) \geq 0$. Lemma 1 therefore highlights a trade-off in the choice of a .

It still remains to bound ρ_1 and ρ_2 away from zero, and that will be done in two different ways for the cases of negative definite Hessians and best expert potentials. Nevertheless, the next technical lemma, which is instrumental in bounding ρ_1 away from zero, still holds in general. Essentially, it says that in the definition of $\rho_1(a)$, it suffices to take the infimum over $\{\|\mathbf{L}\| = a\}$ instead of $\{\|\mathbf{L}\| \geq a\}$.

Lemma 2 *It holds that $\rho_1(a) = \inf_{\|\mathbf{L}\|=a} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\}$.*

Proof Let $\|\mathbf{L}\| \geq a$, and define $\mathbf{L}' = a\mathbf{L}/\|\mathbf{L}\|$, so $\|\mathbf{L}'\| = a$. Since Φ is concave, we have that $\Phi(\mathbf{L}) \geq \Phi(\mathbf{L}') + \nabla \Phi(\mathbf{L}) \cdot (\mathbf{L} - \mathbf{L}')$, and thus

$$\begin{aligned}
\Phi(\mathbf{L}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\} &\geq \Phi(\mathbf{L}') + \nabla \Phi(\mathbf{L}) \cdot (\mathbf{L} - \mathbf{L}') - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\} \\
&\geq \Phi(\mathbf{L}') + \left(\frac{\|\mathbf{L}\|}{a} - 1 \right) \nabla \Phi(\mathbf{L}) \cdot \mathbf{L}' - \frac{\|\mathbf{L}\|}{a} \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}'\} \\
&= \Phi(\mathbf{L}') - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}'\} \\
&\quad + \left(\frac{\|\mathbf{L}\|}{a} - 1 \right) \left(\nabla \Phi(\mathbf{L}) \cdot \mathbf{L}' - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}'\} \right) \\
&\geq \Phi(\mathbf{L}') - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}'\},
\end{aligned}$$

where the last inequality uses the fact that $\nabla \Phi(\mathbf{L}) \cdot \mathbf{L}' - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}'\} \geq 0$. Subtracting $\Phi(\mathbf{0})$, the result follows. $\square$

We now present the main result of this section, assuming we have shown $\rho_1, \rho_2 > 0$. In what follows we write $Q_T = Q_T(\mathbf{l}_1, \dots, \mathbf{l}_T) = \sum_{t=1}^T \|\mathbf{l}_t\|^2$ for the generic quadratic variation w.r.t. the seminorm $\|\cdot\|$ (not to be confused with specific notions of quadratic variation). We denote by $Q > 0$ a given lower bound on Q_T .

Theorem 5 Let $a > 0$ satisfy $\rho_1(a), \rho_2(a) > 0$.

(i) For every $\eta > 0$, $\max_t \{R_{A,t}\} \geq \min\{\frac{\rho_1(a)}{\eta}, \frac{\eta}{2}\rho_2(a)Q\}$, and for $\eta = \sqrt{\frac{2\rho_1(a)}{\rho_2(a)Q}}$,

$$\max_t \{R_{A,t}\} \geq \sqrt{\rho_1(a)\rho_2(a)/2} \cdot \sqrt{Q}.$$

(ii) If for any sequence with quadratic variation $Q_T \leq Q'$ we have $R_{A,T} \leq c\sqrt{Q'}$, then for any such sequence,

$$\max_t \{R_{A,t}\} \geq \frac{\rho_1(a)\rho_2(a)}{2c} \cdot \frac{Q_T}{\sqrt{Q'}}.$$

In particular, if $Q_T = \Theta(Q')$, then $\max_t \{R_{A,t}\} = \Omega(\sqrt{Q_T})$.

Proof (i) By Lemma 1,

$$\max_t \{R_{A,t}\} \geq \min \left\{ \frac{\rho_1(a)}{\eta}, \frac{\eta}{2}\rho_2(a) \sum_{t=1}^T \|\mathbf{l}_t\|^2 \right\} \geq \min \left\{ \frac{\rho_1(a)}{\eta}, \frac{\eta}{2}\rho_2(a)Q \right\}.$$

Picking $\eta = \sqrt{\frac{2\rho_1(a)}{\rho_2(a)Q}}$ implies that $\frac{\rho_1(a)}{\eta} = \frac{\eta}{2}\rho_2(a)Q = \sqrt{\frac{1}{2}Q\rho_1(a)\rho_2(a)}$.

(ii) Given a and η , let $0 < \epsilon < a/\eta$ satisfy that $T = Q'/\epsilon^2$ is an integer. In addition, let $\mathbf{x} \in \mathbb{R}^N$ be such that $\|\mathbf{x}\| = 1$. The loss sequence $\mathbf{l}_t = (-1)^{t+1}\epsilon\mathbf{x}$, for $1 \leq t \leq T$, satisfies that $Q_T = \epsilon^2 T = Q'$ and that $\|\mathbf{l}_t\| = \|(1 - (-1)^t)\epsilon\mathbf{x}/2\| \leq \epsilon < a/\eta$ for every t . Therefore,

$$c\sqrt{Q'} \geq R_{A,T} \geq \frac{\eta}{2}\rho_2(a) \sum_{t=1}^T \|\mathbf{l}_t\|^2 = \frac{\eta}{2}\rho_2(a)Q',$$

where the second inequality is by Lemma 1. This implies that $\eta \leq \frac{2c}{\rho_2(a)\sqrt{Q'}}$.

On the other hand, let $T > \frac{a^2}{\eta^2 Q'}$, define $\epsilon = \sqrt{Q'/T}$, and consider the loss sequence $\mathbf{l}_t = \epsilon\mathbf{x}$. Then $Q_T = \epsilon^2 T = Q'$, and $\|\mathbf{l}_t\| = \epsilon T = \sqrt{Q'T} > a/\eta$. Thus, again using Lemma 1, we have $c\sqrt{Q'} \geq R_{A,T} \geq \rho_1(a)/\eta$, which means that $\eta \geq \frac{\rho_1(a)}{c\sqrt{Q'}}$. Together, we

have that $\frac{\rho_1(a)}{c\sqrt{Q'}} \leq \eta \leq \frac{2c}{\rho_2(a)\sqrt{Q'}}$, implying that $c \geq \sqrt{\rho_1(a)\rho_2(a)/2}$.³ Given any sequence with $Q_T \leq Q'$, we have that $\frac{\rho_1(a)}{\eta} \geq \frac{\rho_1(a)\rho_2(a)\sqrt{Q'}}{2c}$ and $\frac{\eta}{2}\rho_2(a)Q_T \geq \frac{\rho_1(a)\rho_2(a)Q_T}{2c\sqrt{Q'}}$, so by Lemma 1, $\max_t \{R_{A,t}\} \geq \frac{\rho_1(a)\rho_2(a)}{2c} \cdot \frac{Q_T}{\sqrt{Q'}}$, concluding the proof. $\square$

4.1 Potentials with negative definite Hessians

For this case, we pick $\|\cdot\|_2$ as our seminorm. Let $a > 0$ and let $\nabla^2\Phi(\mathbf{L}) \prec 0$ for $\mathbf{L}$ s.t. $\|\mathbf{L}\|_2 \leq a$. In this setting, the infimum in the definitions of $\rho_1(a)$ and $\rho_2(a)$ is equivalent to a minimum, using continuity and the compactness of L^2 balls and spheres.

Lemma 3 *If $\|\cdot\| = \|\cdot\|_2$, then $\rho_1(a) = \min_{\|\mathbf{L}\|=a} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\}$ and $\rho_2(a) = \min_{\|\mathbf{L}\| \leq a, \|\mathbf{I}\|=1} \{-\mathbf{I}^\top \nabla^2\Phi(\mathbf{L})\mathbf{I}\}$.*

Proof By Lemma 2, $\rho_1(a) = \inf_{\|\mathbf{L}\|=a} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\}$. The function $h(\mathbf{L}) = \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}$ is continuous on $\mathbb{R}^N$ (Lemma 10). Therefore, $\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}$ is continuous on the compact set $\{\|\mathbf{L}\| = a\}$ and attains a minimum there. The second statement follows from the fact that $-\mathbf{I}^\top \nabla^2\Phi(\mathbf{L})\mathbf{I}$ is a continuous function on the compact set $\{\|\mathbf{L}\| \leq a\} \times \{\|\mathbf{I}\| = 1\}$, and attains a minimum there. $\square$

By Lemma 3, $\rho_2(a) = \min_{\|\mathbf{L}\| \leq a, \|\mathbf{I}\|=1} \{-\mathbf{I}^\top \nabla^2\Phi(\mathbf{L})\mathbf{I}\} > 0$, where the inequality is true since the Hessians are negative definite, so we are taking the minimum of positive values. In addition, if $\mathbf{L} \neq \mathbf{0}$, then $\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\} > \nabla\Phi(\mathbf{L}) \cdot \mathbf{L} - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\} \geq 0$, since Φ is strictly concave. Thus, again by Lemma 3, $\rho_1(a) = \min_{\|\mathbf{L}\|=a} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\} > 0$. The following statement is an immediate consequence of Theorem 5:

Theorem 6 *If $\nabla^2\Phi(\mathbf{L}) \prec 0$ for every $\mathbf{L}$ s.t. $\|\mathbf{L}\|_2 \leq a$, for some $a > 0$, then*

- (i) *For every $\eta > 0$, it holds that $\max_t \{R_{A,t}\} \geq \min\{\frac{\rho_1(a)}{\eta}, \frac{\eta}{2}\rho_2(a)Q\}$, and for $\eta = \sqrt{\frac{2\rho_1(a)}{\rho_2(a)Q}}$, $\max_t \{R_{A,t}\} \geq \sqrt{\rho_1(a)\rho_2(a)/2} \cdot \sqrt{Q}$.*
- (ii) *If for any sequence with quadratic variation $Q_T \leq Q'$ we have $R_{A,T} \leq c\sqrt{Q'}$, then for any such sequence, $\max_t \{R_{A,t}\} \geq \frac{\rho_1(a)\rho_2(a)}{2c} \cdot \frac{Q_T}{\sqrt{Q'}}$. In particular, if $Q_T = \Theta(Q')$, then $\max_t \{R_{A,t}\} = \Omega(\sqrt{Q_T})$.*

4.2 The best expert setting

In the best expert setting, where $\mathcal{K} = \Delta_N$, potentials can never be strictly concave, let alone have negative definite Hessians. To see that, let $\mathbf{L} \in \mathbb{R}^N$, $c \in \mathbb{R}$, and define $\mathbf{L}' = \mathbf{L} + c \cdot \mathbf{1}$, where $\mathbf{1}$ is the all-one vector. We will say that $\mathbf{L}'$ is a *uniform translation* of $\mathbf{L}$. Then

$$c = \nabla\Phi(\mathbf{L}) \cdot (\mathbf{L}' - \mathbf{L}) \geq \Phi(\mathbf{L}') - \Phi(\mathbf{L}) \geq \nabla\Phi(\mathbf{L}') \cdot (\mathbf{L}' - \mathbf{L}) = c,$$

where we use the concavity of Φ , the fact that $\nabla\Phi$ is a probability vector, and the fact that $\mathbf{L}' - \mathbf{L} = c \cdot \mathbf{1}$. For a strictly concave Φ , the above inequalities would be strict if $c \neq 0$, but instead, they are equalities. Thus, the conditions for strict concavity are not fulfilled at any point $\mathbf{L}$.

We will replace the negative definite assumption with the assumption that for every $i \neq j$, $\frac{\partial^2\Phi}{\partial L_i \partial L_j} > 0$. This condition is natural for regret minimization algorithms, because

³ Note that this means we cannot guarantee a regret upper bound of $c\sqrt{Q'}$ for $c < \sqrt{\rho_1(a)\rho_2(a)/2}$.

$\frac{\partial^2 \Phi(\mathbf{L})}{\partial L_i \partial L_j} = \frac{\partial p_i(\mathbf{L})}{\partial L_j}$, where p_i is the weight of expert i . Thus, we simply require that an increase in the cumulative loss of expert j results in an increase in the weight of every other expert (and hence a decrease in its own weight). A direct implication of this assumption is that $\frac{\partial \Phi(\mathbf{L})}{\partial L_i} > 0$ for every i and $\mathbf{L}$. To see that, observe that $p_i(\mathbf{L}) = 1 - \sum_{j \neq i} p_j(\mathbf{L})$, so $\frac{\partial p_i(\mathbf{L})}{\partial L_i} = -\sum_{j \neq i} \frac{\partial p_j(\mathbf{L})}{\partial L_i} < 0$. Since $p_i(\mathbf{L}) \geq 0$ and it is strictly decreasing in L_i , it follows that $p_i(\mathbf{L}) > 0$, or $\frac{\partial \Phi(\mathbf{L})}{\partial L_i} > 0$.

Using the above assumption we proceed to bound ρ_1 and ρ_2 away from zero. We first state some general properties of best expert potentials.

Lemma 4 (i) Every row and column of $\nabla^2 \Phi$ sum up to zero. (ii) $\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}$ is invariant w.r.t. a uniform translation of $\mathbf{L}$. (iii) $\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{1}$ is invariant w.r.t. a uniform translation of either $\mathbf{1}$ or $\mathbf{L}$.

Proof (i) Defining $h(\mathbf{L}) = \sum_{i=1}^N \frac{\partial \Phi(\mathbf{L})}{\partial L_i}$, we have that $h \equiv 1$, and thus $0 = \frac{\partial h(\mathbf{L})}{\partial L_j} = \sum_{i=1}^N \frac{\partial^2 \Phi(\mathbf{L})}{\partial L_i \partial L_j}$ for every j .

(ii) Let $\mathbf{L}' = \mathbf{L} + c \cdot \mathbf{1}$. As already seen, $\Phi(\mathbf{L}') - \Phi(\mathbf{L}) = c$, and we also have that

$$\min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}'\} = \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L} + c \cdot \mathbf{u} \cdot \mathbf{1}\} = \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\} + c.$$

Therefore,

$$\Phi(\mathbf{L}') - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}'\} = \Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}.$$

(iii) Note that in this context, $\mathbf{1}, \mathbf{L} \in \mathbb{R}^N$ are unrelated vectors.

$$\begin{aligned} (\mathbf{1} + c\mathbf{1})^\top \nabla^2 \Phi(\mathbf{L})(\mathbf{1} + c\mathbf{1}) &= \mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{1} + c\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L})(\mathbf{1} + c\mathbf{1}) + \mathbf{1}^\top \nabla^2 \Phi(\mathbf{L})(c\mathbf{1}) \\ &= \mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{1}, \end{aligned}$$

where we have that $\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) = \mathbf{0}^\top$ and $\nabla^2 \Phi(\mathbf{L}) \mathbf{1} = \mathbf{0}$ by (i). To show the invariance for $\mathbf{L}$, recall that for $\mathbf{L}' = \mathbf{L} + c \cdot \mathbf{1}$, $\Phi(\mathbf{L}') - \Phi(\mathbf{L}) = c$, and thus, $\nabla^2 \Phi(\mathbf{L}') = \nabla^2 \Phi(\mathbf{L})$. $\square$

We now consider $\rho_1(a)$ and $\rho_2(a)$ where we use the seminorm $\|\mathbf{v}\| = \delta(\mathbf{v})$. (The fact that $\delta(\mathbf{v}) = \max_i \{v_i\} - \min_i \{v_i\}$ is a seminorm is proved in Lemma 11 in Appendix 1.) Under this seminorm, $\sum_{t=1}^T \|\mathbf{l}_t\|^2$ becomes q_T , the relative quadratic variation. Note that $\delta(\mathbf{v})$ is invariant to uniform translation. In particular, for every $\mathbf{v} \in \mathbb{R}^N$ we may consider its “normalized” version, $\hat{\mathbf{v}} = \mathbf{v} - \min_i \{v_i\} \cdot \mathbf{1}$. We have that $\delta(\hat{\mathbf{v}}) = \delta(\mathbf{v})$, $\hat{\mathbf{v}} \in [0, \delta(\mathbf{v})]^N$, and there exist entries i and j s.t. $\hat{v}_i = 0$ and $\hat{v}_j = \delta(\mathbf{v})$. Writing $\mathcal{N}(a)$ for the set of normalized vectors with seminorm a , we thus have that $\mathcal{N}(a) = \{\mathbf{v} \in [0, a]^N : \exists i, j \text{ s.t. } v_i = a, v_j = 0\}$. The set $\mathcal{N}(a)$ is bounded and also closed, as a finite union and intersection of closed sets.

Using invariance to uniform translation, we can now show that the infima in the expressions for ρ_1 and ρ_2 may be taken over compact sets, and thus be replaced with minima. Using the requirement that $\frac{\partial^2 \Phi}{\partial L_i \partial L_j} > 0$ for every $i \neq j$, we can then show that the expressions inside the minima are positive. This is summarized in the next lemma.

Lemma 5 For the best expert setting, it holds that $\rho_1(a) = \min_{\mathbf{L} \in \mathcal{N}(a)} \{\Phi(\mathbf{L})\} - \Phi(\mathbf{0}) > 0$ and $\rho_2(a) = \min_{\mathbf{L} \in [0, a]^N, \mathbf{1} \in \mathcal{N}(1)} \{-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{1}\} > 0$.

Proof We start with $\rho_1(a)$. By Lemma 2, we have that $\rho_1(a) = \inf_{\|\mathbf{L}\|=a} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\}$. By part (ii) of Lemma 4, $\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}$ is invariant to uniform translation, so $\rho_1(a) = \inf_{\mathbf{L} \in \mathcal{N}(a)} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\}$. For every normalized vector $\mathbf{L}$, $\min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\} = 0$, by placing all the weight of $\mathbf{u}$ on the zero entry of $\mathbf{L}$. In addition, as already observed, $\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}$ is continuous, and by the compactness of $\mathcal{N}(a)$, $\rho_1(a) = \min_{\mathbf{L} \in \mathcal{N}(a)} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0}) - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}\} = \min_{\mathbf{L} \in \mathcal{N}(a)} \{\Phi(\mathbf{L}) - \Phi(\mathbf{0})\}$. Now, $\Phi(\mathbf{L}) - \Phi(\mathbf{0}) \geq \nabla \Phi(\mathbf{L}) \cdot \mathbf{L} = \sum_{i=1}^N p_i(\mathbf{L}) L_i$ by the concavity of Φ . Recall that $p_i(\mathbf{L}) > 0$, and for $\mathbf{L} \in \mathcal{N}(a)$, we have that $L_i \geq 0$ for every i , and there is at least one index with $L_i > 0$. Thus, $\Phi(\mathbf{L}) - \Phi(\mathbf{0}) > 0$, proving the statement for ρ_1 .

We now move on to ρ_2 . By Lemma 4, $-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}$ is invariant to a uniform translation of either $\mathbf{l}$ or $\mathbf{L}$. Therefore, $\rho_2(a) = \inf_{\mathbf{L} \in [0, a]^N, \mathbf{l} \in \mathcal{N}(1)} \{-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}\}$. Since $-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}$ is continuous on the compact set $[0, a]^N \times \mathcal{N}(1)$, $\rho_2(a) = \min_{\mathbf{L} \in [0, a]^N, \mathbf{l} \in \mathcal{N}(1)} \{-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}\}$. To show that $\rho_2(a) > 0$, it therefore suffices to show that $-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l} > 0$ for every $\mathbf{L} \in [0, a]^N$ and $\mathbf{l} \in \mathcal{N}(1)$. Fix $\mathbf{L}$, and denote $A = -\nabla^2 \Phi(\mathbf{L})$. For any $\mathbf{l} \in \mathcal{N}(1)$, one entry is 1, another is 0, and the rest are in $[0, 1]$. If $N = 2$, we have that $-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}$ is either $-\nabla^2 \Phi(\mathbf{L})_{1,1}$ or $-\nabla^2 \Phi(\mathbf{L})_{2,2}$, and in both cases positive. We next assume $N \geq 3$ and minimize $f(\mathbf{l}) = \mathbf{l}^\top A \mathbf{l} = \sum_{i,j} a_{i,j} l_i l_j$ for $l_k \in [0, 1]$, $k \geq 3$, where $l_1 = 1, l_2 = 0$. For every k ,

$$\frac{\partial f}{\partial l_k} = 2a_{k,k}l_k + 2 \sum_{i \neq k} a_{i,k}l_i = 2 \sum_{i \neq k} a_{i,k}(l_i - l_k),$$

where the last equality is by part (i) of Lemma 4. Note that $a_{i,k} < 0$ if $i \neq k$ and $a_{k,k} > 0$. Thus, for $l_k = 0$, $\frac{\partial f}{\partial l_k} < 0$, and for $l_k = 1$, $\frac{\partial f}{\partial l_k} > 0$, so by the linearity of $\frac{\partial f}{\partial l_k}$, there is $l_k^* \in [0, 1]$ satisfying $\frac{\partial f}{\partial l_k} = 0$, and it is a minimizing choice for f . Thus, at a minimum point, it holds that $\frac{\partial f}{\partial l_k} = 0$, or $\sum_{i \neq k} a_{i,k}l_i = -a_{k,k}l_k$ for every $k \geq 3$. We have that

$$\begin{aligned} f(\mathbf{l}) &= \sum_{i=1}^N a_{i,i}l_i^2 + \sum_{i=1}^N \sum_{j \neq i} a_{i,j}l_i l_j = \sum_{i=1}^N a_{i,i}l_i^2 + \sum_{i=1}^2 l_i \sum_{j \neq i} a_{i,j}l_j + \sum_{i=3}^N l_i \sum_{j \neq i} a_{i,j}l_j \\ &= \sum_{i=1}^N a_{i,i}l_i^2 + \sum_{j \neq 1} a_{1,j}l_j + \sum_{i=3}^N l_i(-a_{i,i}l_i) = \sum_{i=1}^N a_{i,i}l_i^2 + \sum_{j \neq 1} a_{1,j}l_j - \sum_{i=3}^N a_{i,i}l_i^2 \\ &= \sum_{i=1}^2 a_{i,i}l_i^2 + \sum_{j \neq 1} a_{1,j}l_j = a_{1,1} + \sum_{j \neq 1} a_{1,j}l_j = - \sum_{j \neq 1} a_{1,j} + \sum_{j \neq 1} a_{1,j}l_j \\ &= \sum_{j \neq 1} a_{1,j}(l_j - 1) \geq a_{1,2}(l_2 - 1) = -a_{1,2}. \end{aligned}$$

Since our choice of $l_1 = 1$ and $l_2 = 0$ was arbitrary, we have that $f(\mathbf{l}) \geq \min_{i \neq j} \{-a_{i,j}\} > 0$. Thus, $-\mathbf{1}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l} > 0$, completing the proof. $\square$

We can now apply Theorem 5 to the best expert setting.

Theorem 7 If $\frac{\partial^2 \Phi}{\partial L_i \partial L_j} > 0$ for every $i \neq j$ and every $\mathbf{L}$ s.t. $\delta(\mathbf{L}) \leq a$, then

- (i) For every $\eta > 0$, it holds that $\max_t \{R_{A,t}\} \geq \min\{\frac{\rho_1(a)}{\eta}, \frac{\eta}{2} \rho_2(a)q\}$, and for $\eta = \sqrt{\frac{2\rho_1(a)}{\rho_2(a)q}}$, $\max_t \{R_{A,t}\} \geq \sqrt{\rho_1(a)\rho_2(a)/2} \cdot \sqrt{q}$.

- (ii) If for any sequence with relative quadratic variation $q_T \leq q'$ we have $R_{A,T} \leq c\sqrt{q'}$, then for any such sequence, $\max_t \{R_{A,t}\} \geq \frac{\rho_1(a)\rho_2(a)}{2c} \cdot \frac{q_T}{\sqrt{q'}}$. In particular, if $q_T = \Theta(q')$, then $\max_t \{R_{A,t}\} = \Omega(\sqrt{q_T})$.

5 Application to specific regret minimization algorithms

5.1 Online Gradient Descent with linear costs

In this subsection, we deal with the Lazy Projection variant of the OGD algorithm (Zinkevich 2003) with a fixed learning rate η and linear costs. In this setting, for each t , OGD selects a weight vector according to the rule $\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{K}} \{\|\mathbf{x} + \eta \mathbf{L}_t\|_2\}$, where $\mathcal{K} \subset \mathbb{R}^N$ is compact and convex. As observed in Hazan and Kale (2010) and Hazan (2011), this algorithm is equivalent to $RFTL(\eta, \mathcal{R})$, where $\mathcal{R}(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$, namely, setting $\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{K}} \{\mathbf{x} \cdot \mathbf{L}_t + (1/2\eta)\|\mathbf{x}\|_2^2\}$. In what follows we will make the assumption that $\mathcal{K} \supseteq B(\mathbf{0}, a)$, where $B(\mathbf{0}, a)$ is the closed ball with radius a centered at $\mathbf{0}$, for some $a > 0$.

Note that solving the above minimization problem without the restriction $\mathbf{x} \in \mathcal{K}$ yields $\mathbf{x}'_{t+1} = -\eta \mathbf{L}_t$. However, if $\|\mathbf{L}_t\|_2 \leq a/\eta$, then $\mathbf{x}'_{t+1} \in \mathcal{K}$, and then, in fact, $\mathbf{x}_{t+1} = -\eta \mathbf{L}_t$. By Theorem 3,

$$\begin{aligned}\Phi_\eta(\mathbf{L}_t) &= \mathbf{x}_{t+1} \cdot \mathbf{L}_t + (1/\eta)\mathcal{R}(\mathbf{x}_{t+1}) = -\eta \mathbf{L}_t \cdot \mathbf{L}_t + (1/2\eta)\|\mathbf{L}_t\|_2^2 \\ &= -(\eta/2)\|\mathbf{L}_t\|_2^2.\end{aligned}$$

Thus, if $\|\mathbf{L}\|_2 \leq a$, then $\Phi(\mathbf{L}) = -(1/2)\|\mathbf{L}\|_2^2$ and also $\nabla^2 \Phi(\mathbf{L}) = -I$, where I is the identity matrix. By Lemma 3,

$$\rho_1(a) = \min_{\|\mathbf{L}\|_2=a} \left\{ -\frac{1}{2}\|\mathbf{L}\|_2^2 - \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\} \right\} \geq \min_{\|\mathbf{L}\|_2=a} \left\{ -\frac{1}{2}\|\mathbf{L}\|_2^2 - (-\mathbf{L}) \cdot \mathbf{L} \right\} = \frac{1}{2}a^2,$$

where we used the fact that $-\mathbf{L} \in \mathcal{K}$ if $\|\mathbf{L}\|_2 = a$. In addition, by Lemma 3,

$$\rho_2(a) = \min_{\|\mathbf{L}\|_2 \leq a, \|\mathbf{I}\|_2=1} \{-\mathbf{I}^\top (-I) \mathbf{I}\} = 1.$$

By Theorem 6, we have that $\max_t \{R_{A,t}\} \geq \min\{a^2/(2\eta), (\eta/2)Q\}$, and for $\eta = \frac{a}{\sqrt{Q}}$, $\max_t \{R_{A,t}\} \geq \frac{a}{2}\sqrt{Q}$.

5.2 Randomized Weighted Majority

RWM is the most notable regret minimization algorithm for the expert setting. We have $\mathcal{K} = \Delta_N$, and the algorithm gives a weight $p_{i,t+1} = \frac{p_{i,0}e^{-\eta L_{i,t}}}{\sum_{j=1}^N p_{j,0}e^{-\eta L_{j,t}}}$ to expert i at time $t+1$, where the initial weights $p_{i,0}$ and the learning rate η are parameters.

It is easy to see that for the potential $\Phi_\eta(\mathbf{L}) = -(1/\eta) \ln(\sum_{i=1}^N p_{i,0}e^{-\eta L_i})$, we have that $\mathbf{p} = (p_1, \dots, p_N) = \nabla \Phi_\eta(\mathbf{L})$. The Hessian $\nabla^2 \Phi_\eta$ has the following simple form:

Lemma 6 Let $\mathbf{L} \in \mathbb{R}^N$ and denote $\mathbf{p} = \nabla \Phi_\eta(\mathbf{L})$. Then $\nabla^2 \Phi_\eta(\mathbf{L}) = \eta \cdot (\mathbf{p}\mathbf{p}^\top - \text{diag}(\mathbf{p})) \leq 0$, where $\text{diag}(\mathbf{p})$ is the diagonal matrix with $\mathbf{p}$ as its diagonal.

The proof is given in Appendix 2. We will assume $p_{1,0} = \dots = p_{N,0} = 1/N$, and write $RWM(\eta)$ for RWM with parameters η and the uniform distribution. Thus, $\Phi(\mathbf{L}) =$

$-\ln((1/N) \sum_{i=1}^N e^{-L_i})$, and we have by Lemma 6 that $\frac{\partial^2 \Phi}{\partial L_i \partial L_j} > 0$ for every $i \neq j$. Therefore, by Lemma 5, $\rho_1, \rho_2 > 0$. We now need to calculate ρ_1 and ρ_2 . This is straightforward in the case of ρ_1 , but for ρ_2 we give the value only for $N = 2$.

Lemma 7 For any $N \geq 2$, $\rho_1(a) = \ln \frac{N}{N-1+e^{-a}}$. For $N = 2$, it holds that $\rho_2(a) = (e^{a/2} + e^{-a/2})^{-2}$.

Proof By Lemma 5,

$$\begin{aligned} \rho_1(a) &= \min_{\mathbf{L} \in \mathcal{N}(a)} \{\Phi(\mathbf{L})\} - \Phi(\mathbf{0}) = \min_{\mathbf{L} \in \mathcal{N}(a)} \{\Phi(\mathbf{L})\} = -\ln \max_{\mathbf{L} \in \mathcal{N}(a)} \left\{ \frac{1}{N} \sum_{i=1}^N e^{-L_i} \right\} \\ &= -\ln \left(\frac{1}{N} ((N-1)e^{-0} + e^{-a}) \right) = \ln \frac{N}{N-1+e^{-a}}. \end{aligned}$$

In addition, $\rho_2(a) = \min_{\mathbf{L} \in [0,a]^N, \mathbf{I} \in \mathcal{N}(1)} \{-\mathbf{I}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{I}\} > 0$. For $N = 2$, we have $\mathcal{N}(1) = \{(1, 0), (0, 1)\}$. Therefore, $-\mathbf{I}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{I}$ is either $-\nabla^2 \Phi(\mathbf{L})_{1,1}$ or $-\nabla^2 \Phi(\mathbf{L})_{2,2}$. Denoting $\mathbf{p} = \mathbf{p}(\mathbf{L}) = \nabla \Phi(\mathbf{L})$, we have that $-\nabla^2 \Phi(\mathbf{L})_{1,1} = p_1 - p_1^2$ and $-\nabla^2 \Phi(\mathbf{L})_{2,2} = p_2 - p_2^2$ by Lemma 6. Since $p_1 = 1 - p_2$, we have $p_1 - p_1^2 = p_2 - p_2^2$. Now,

$$\begin{aligned} p_1(1 - p_1) &= \frac{e^{-L_1}}{e^{-L_1} + e^{-L_2}} \cdot \frac{e^{-L_2}}{e^{-L_1} + e^{-L_2}} = e^{-(L_1+L_2)} (e^{-L_1} + e^{-L_2})^{-2} \\ &= (e^{(L_2-L_1)/2} + e^{(L_1-L_2)/2})^{-2}. \end{aligned}$$

The function $(e^x + e^{-x})^{-2}$ is decreasing for $x \geq 0$, so the smallest value is attained for $|L_1 - L_2| = a$. Thus, for $N = 2$, $\rho_2(a) = (e^{a/2} + e^{-a/2})^{-2}$. $\square$

Picking $a = 1.2$, we have by Theorem 7 that

Theorem 8 For $N = 2$, there exists η s.t.

$$\max_t \{R_{RWM(\eta),t}\} \geq \sqrt{\rho_1(a)\rho_2(a)q/2} \geq 0.195\sqrt{q}.$$

For a general N we may still lower bound $\max_t \{R_{RWM(\eta),t}\}$ by providing a lower bound for $\rho_2(a)$. We use the fact that the term $-\mathbf{I}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{I}$, which is minimized in the definition of $\rho_2(a)$, may be interpreted as the variance of a certain discrete and bounded random variable. Thus, the key element that will be used is a general lower bound on the variance of such variables. This bound, possibly of separate interest, is given in the following lemma, whose proof can be found in Appendix 2.

Lemma 8 Let $x_1 < \dots < x_N \in \mathbb{R}$, $0 < p_1, \dots, p_N < 1$, and $\sum_{k=1}^N p_k = 1$, and let X be a random variable that obtains the value x_k with probability p_k , for every $1 \leq k \leq N$. Then for every $1 \leq i < j \leq N$, $\text{Var}(X) \geq \frac{p_i p_j}{p_i + p_j} \cdot (x_i - x_j)^2$, and equality is attained iff $N = 2$, or $N = 3$ with $i = 1, j = 3$, and $x_2 = \frac{p_1 x_1 + p_3 x_3}{p_1 + p_3}$. In particular, $\text{Var}(X) \geq \frac{p_1 p_N}{p_1 + p_N} \cdot (x_N - x_1)^2 \geq \frac{1}{2} \min_k \{p_k\} (x_N - x_1)^2$.

We comment that for the special case $p_1 = \dots = p_N = 1/N$, Lemma 8 yields the inequality $\text{Var}(X) \geq (x_N - x_1)^2 / (2N)$, sometimes referred to as the Szökefalvi Nagy inequality (von Szökefalvi Nagy 1918).

We now proceed to lower bound $\rho_2(a)$ for any N .

Lemma 9 For any $N \geq 2$, it holds that $\rho_2(a) \geq \frac{1}{2(N-1)e^a + 2}$.

Proof Recall that $\rho_2(a) = \min_{\mathbf{L} \in [0, a]^N, \mathbf{l} \in \mathcal{N}(1)} \{-\mathbf{l}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}\}$. By Lemma 6, $\nabla^2 \Phi(\mathbf{L}) = \mathbf{p}\mathbf{p}^\top - \text{diag}(\mathbf{p})$, where $\mathbf{p} = \mathbf{p}(\mathbf{L}) = \nabla \Phi(\mathbf{L})$, and thus,

$$-\mathbf{l}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l} = \mathbf{l}^\top \text{diag}(\mathbf{p}) \mathbf{l} - \mathbf{l}^\top \mathbf{p}\mathbf{p}^\top \mathbf{l} = \sum_{i=1}^N p_i l_i^2 - \left(\sum_{i=1}^N p_i l_i \right)^2.$$

Assume for now that $l_1, \dots, l_N$ are distinct. Then $-\mathbf{l}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l} = \text{Var}(X)$, where X is a random variable obtaining the value l_i with probability p_i , $i = 1, \dots, N$. By Lemma 8, $\text{Var}(X) \geq \frac{1}{2} \min_k \{p_k\} (\max_i \{l_i\} - \min_i \{l_i\})^2$. For $\mathbf{l} \in \mathcal{N}(1)$ it holds that $\max_i \{l_i\} = 1$ and $\min_i \{l_i\} = 0$. In addition, for $\mathbf{L} \in [0, a]^N$, the smallest value for $\min_k \{p_k\}$ is obtained when one entry of $\mathbf{L}$ is a and the rest are 0, yielding that $\min_k \{p_k\} \geq \frac{e^{-a}}{N-1+e^{-a}}$. Together we have that $-\mathbf{l}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l} \geq \frac{1}{2} \cdot \frac{e^{-a}}{N-1+e^{-a}} = \frac{1}{2(N-1)e^a+2}$. Since the expression $-\mathbf{l}^\top \nabla^2 \Phi(\mathbf{L}) \mathbf{l}$ is continuous in $\mathbf{l}$, this inequality holds even if the assumption that the entries of $\mathbf{l}$ are distinct numbers is dropped. The result follows. $\square$

Theorem 9 For any $0 < \alpha < 1/N$ and $\eta > 0$, it holds that

$$\max_t \{R_{RWM(\eta),t}\} \geq \min \left\{ \frac{1}{\eta} \ln \frac{N(1-\alpha)}{N-1}, \frac{1}{4} \eta \alpha q \right\},$$

and for $\alpha = 1/(2N)$ and $\eta = \sqrt{(8N/q) \ln \frac{2N-1}{2N-2}}$,

$$\max_t \{R_{RWM(\eta),t}\} \geq \frac{\sqrt{q}}{4N} \cdot \sqrt{\frac{2N}{2N-1}} = \Omega(\sqrt{q}/N).$$

Proof By Theorem 7, for every $\eta > 0$, it holds that

$$\max_t \{R_{RWM(\eta),t}\} \geq \min \left\{ \frac{\rho_1(a)}{\eta}, \frac{\eta}{2} \rho_2(a) q \right\}.$$

We fix a , observing that any $a \in (0, \infty)$ may be used, and denote $\alpha = \frac{1}{(N-1)e^a+1}$. Note that α may thus obtain any value in $(0, 1/N)$. By Lemma 9, we have that $\rho_2(a) \geq \alpha/2$. By Lemma 7, $\rho_1(a) = \ln \frac{N}{N-1+e^{-a}}$, and since

$$\frac{N(1-\alpha)}{N-1} = \frac{N}{N-1} \cdot \frac{(N-1)e^a}{(N-1)e^a+1} = \frac{Ne^a}{(N-1)e^a+1} = \frac{N}{N-1+e^{-a}},$$

we have that $\rho_1(a) = \ln \frac{N(1-\alpha)}{N-1}$. Thus, $\max_t \{R_{RWM(\eta),t}\} \geq \min \left\{ \frac{1}{\eta} \ln \frac{N(1-\alpha)}{N-1}, \frac{1}{4} \eta \alpha q \right\}$. We may maximize the r.h.s. by picking $\eta = \sqrt{\frac{4}{\alpha q} \ln \frac{N(1-\alpha)}{N-1}}$, yielding $\max_t \{R_{RWM(\eta),t}\} \geq \sqrt{\frac{1}{4} \alpha q \ln \frac{N(1-\alpha)}{N-1}}$. We may further pick a s.t. $\alpha = 1/(2N)$ and obtain

$$\begin{aligned} \max_t \{R_{RWM(\eta),t}\} &\geq \sqrt{\frac{q}{8N} \ln \frac{N-1/2}{N-1}} = \sqrt{\frac{q}{8N} \ln \left(1 + \frac{1}{2N-2} \right)} \\ &\geq \sqrt{\frac{q}{8N} \cdot \frac{\frac{1}{2N-2}}{1 + \frac{1}{2N-2}}} = \sqrt{\frac{q}{8N} \cdot \frac{1}{2N-1}} \\ &= \frac{\sqrt{q}}{4N} \cdot \sqrt{\frac{2N}{2N-1}}, \end{aligned}$$

where the second inequality uses the fact that $\ln(1+x) \geq \frac{x}{x+1}$ for $x \in (-1, \infty)$. $\square$

6 Beyond quadratic variation

Up until now we have considered only quadratic variation as a measure for the smoothness of the loss sequence. In this section we will extend our results to more general measures, which include other notions of variation discussed in the literature (Hazan and Kale 2010; Chiang et al. 2012).

In what follows, we will consider some general quantity V_t that measures the variation of the loss sequence at time t , and denote by V' a known upper bound on V_T and by $V > 0$ a known lower bound on V_T . We will further assume that for every loss sequence $\mathbf{l}_1, \dots, \mathbf{l}_T \in \mathbb{R}^N$ it holds that $Q_T \geq \beta V_T$ for some $\beta = \beta(N) > 0$, where Q_T is the generic quadratic variation w.r.t. some seminorm (not to be confused with the definition for a specific seminorm). This restriction will still allow us to apply our results to the most interesting notions of variation, as will be shown later.

We now proceed to extend our original result, given in Theorem 5, to accommodate more general variation. Recall that we assume that $a > 0$ satisfies $\rho_1(a), \rho_2(a) > 0$, and let $\eta > 0$. The first part of Theorem 5 states that for a concave potential algorithm A , it holds that $\max_t \{R_{A,t}\} \geq \min\{\frac{\rho_1(a)}{\eta}, \frac{\eta}{2}\rho_2(a)Q\}$, where Q is a known lower bound on Q_T . In our present context we may clearly set $Q = \beta V$. It immediately follows that for $\eta = \sqrt{\frac{2\rho_1(a)}{\rho_2(a)\beta V}}$, it holds that $\max_t \{R_{A,t}\} \geq \sqrt{\rho_1(a)\rho_2(a)\beta/2} \cdot \sqrt{V}$.

The proof of the second part of the theorem in the context of general variation is also quite similar to the original one and will be given for completeness. Here we assume that for any loss sequence satisfying $V_T \leq V'$, the algorithm A further guarantees $R_{A,T} \leq c\sqrt{V'}$ for some $c > 0$.

For given a and η , there exists $0 < \epsilon < a/\eta$ s.t. $T = \beta V'/\epsilon^2$ is an integer. In addition, let $\mathbf{x} \in \mathbb{R}^N$ be such that $\|\mathbf{x}\| = 1$. For the loss sequence $\mathbf{l}_t = (-1)^{t+1}\epsilon\mathbf{x}$, $1 \leq t \leq T$, it holds that $\beta V_T \leq Q_T = \epsilon^2 T = \beta V'$ and that $\|\mathbf{l}_t\| = \|(1 - (-1)^t)\epsilon\mathbf{x}/2\| \leq \epsilon < a/\eta$ for every t . Therefore,

$$c\sqrt{V'} \geq R_{A,T} \geq \frac{\eta}{2}\rho_2(a)Q_T = \frac{\eta}{2}\rho_2(a)\beta V',$$

where the second inequality is by Lemma 1. This implies that $\eta \leq \frac{2c}{\rho_2(a)\beta\sqrt{V'}}$.

Next, let $T > \frac{a^2}{\eta^2\beta V'}$, define $\epsilon = \sqrt{\beta V'/T}$, and consider the loss sequence $\mathbf{l}_t = \epsilon\mathbf{x}$. Then $\beta V_T \leq Q_T = \epsilon^2 T = \beta V'$, and $\|\mathbf{l}_T\| = \epsilon T = \sqrt{\beta V'T} > a/\eta$. Therefore, again using Lemma 1, we have $c\sqrt{V'} \geq R_{A,T} \geq \rho_1(a)/\eta$, and thus $\eta \geq \frac{\rho_1(a)}{c\sqrt{V'}}$. Together, we obtain that $\frac{\rho_1(a)}{c\sqrt{V'}} \leq \eta \leq \frac{2c}{\rho_2(a)\beta\sqrt{V'}}$, implying that $c \geq \sqrt{\rho_1(a)\rho_2(a)\beta/2}$. (Note that it is therefore impossible to guarantee a regret upper bound of $c\sqrt{V'}$ for $c < \sqrt{\rho_1(a)\rho_2(a)\beta/2}$.) For any sequence with $V_T \leq V'$, we have that $\frac{\rho_1(a)}{\eta} \geq \frac{\rho_1(a)\rho_2(a)\beta\sqrt{V'}}{2c}$ and

$$\frac{\eta}{2}\rho_2(a)Q_T \geq \frac{\rho_1(a)\rho_2(a)Q_T}{2c\sqrt{V'}} \geq \frac{\rho_1(a)\rho_2(a)\beta V_T}{2c\sqrt{V'}},$$

so by Lemma 1, $\max_t \{R_{A,t}\} \geq \frac{\rho_1(a)\rho_2(a)\beta}{2c} \cdot \frac{V_T}{\sqrt{V'}}$. We may therefore generalize Theorem 5 as follows:

Theorem 10 Let $a > 0$ satisfy $\rho_1(a), \rho_2(a) > 0$.

(i) For every $\eta > 0$, $\max_t \{R_{A,t}\} \geq \min\{\frac{\rho_1(a)}{\eta}, \frac{\eta}{2}\rho_2(a)\beta V\}$, and for $\eta = \sqrt{\frac{2\rho_1(a)}{\rho_2(a)\beta V}}$,

$$\max_t \{R_{A,t}\} \geq \sqrt{\rho_1(a)\rho_2(a)\beta/2} \cdot \sqrt{V}.$$

(ii) If for any sequence with variation $V_T \leq V'$ we have $R_{A,T} \leq c\sqrt{V'}$, then for any such sequence,

$$\max_t \{R_{A,t}\} \geq \frac{\rho_1(a)\rho_2(a)\beta}{2c} \cdot \frac{V_T}{\sqrt{V'}}.$$

In particular, if $V_T = \Theta(V')$, then $\max_t \{R_{A,t}\} = \Omega(\sqrt{V_T})$.

6.1 Application to specific notions of variation

The work of [Hazan and Kale \(2010\)](#) defines several types of variation for a loss sequence. In the context of online linear optimization they define $\text{Var}_T = \sum_{t=1}^T \|\mathbf{l}_t - \boldsymbol{\mu}_T\|_2^2$, where $\boldsymbol{\mu}_T = \frac{1}{T} \sum_{t=1}^T \mathbf{l}_t$. In addition, for the best expert setting they define $\text{Var}_T^\infty = \max_{1 \leq i \leq N} \{\sum_{t=1}^T |l_{i,t} - \mu_{i,T}|^2\}$, where $\boldsymbol{\mu}_T$ is as before; they further define $\text{Var}_T^{\max} = \max_{t \leq T} \{\text{Var}_t(m(t))\}$, where $m(t)$ is the best expert up to time t and $\text{Var}_t(i)$ is the variation of the losses of expert i up to time t . It is clear that $\text{Var}_T^{\max} \leq \text{Var}_T^\infty \leq \text{Var}_T$, and furthermore, $\text{Var}_T \leq Q_T$, where the seminorm is taken to be the Euclidean norm. Thus, in the context of Theorem 10 we may choose $\beta = 1$ for all these notions of variation.

The work of [Chiang et al. \(2012\)](#) considers the L_p -deviation of a loss sequence, defined in the context of online linear optimization by $D_p = \sum_{t=1}^T \|\mathbf{l}_t - \mathbf{l}_{t-1}\|_p^2$, where $\mathbf{l}_0 = \mathbf{0}$. Since for every $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N$ we have that $\|\mathbf{x} - \mathbf{y}\|_p^2 \leq (\|\mathbf{x}\|_p + \|\mathbf{y}\|_p)^2 \leq 2(\|\mathbf{x}\|_p^2 + \|\mathbf{y}\|_p^2)$, it follows that $D_p \leq 4 \sum_{t=1}^T \|\mathbf{l}_t\|_p^2$. Thus, we have that for the L_p -deviation we may take $\beta = 4$ (where Q_T is defined w.r.t. the L_p norm).

It is important to note, however, that in order to satisfy the conditions of Theorem 10, one requires a relevant concave potential algorithm. In particular, for the second part of the theorem one needs an algorithm with a suitable regret minimization guarantee. For Var_T such a choice exists in the form of OGD. As shown by [Hazan and Kale \(2010\)](#), if Var_T is upper bounded by $V' \geq 12$ and (w.l.o.g.) $\mathcal{K} \subseteq B(\mathbf{0}, 1)$, then for bounded single-period losses (w.l.o.g. $\|\mathbf{l}_t\|_2 \leq 1$ for every t), the regret of OGD is upper bounded by $15\sqrt{V'}$. In addition, if $\mathcal{K} \supseteq B(\mathbf{0}, a)$ for some $a > 0$, then $\rho_1(a), \rho_2(a) > 0$ (as already shown in Sect. 5.1) and Theorem 10 may be meaningfully applied.⁴

A second algorithm, Variation MW, given in [Hazan and Kale \(2010\)](#) specifically for the best expert setting, is not a concave potential algorithm. This is evident from the fact that its weights have an asymmetric dependence on single-period losses from different times. (This cannot be the case when the weight is a function of the cumulative loss.) The same is true for the algorithm of [Chiang et al. \(2012\)](#) when applied to the best expert setting. It is not immediately clear if there are interesting cases for which their algorithm is equivalent to a concave potential algorithm.

We summarize the results of this subsection in the next two corollaries.

Corollary 2 Theorem 10 may be applied to the following types of variation:

(i) To $\text{Var}_T^{\max}$, Var_T^∞ , and Var_T , with $Q_T = \sum_{t=1}^T \|\mathbf{l}_t\|_2^2$ and $\beta = 1$.

⁴ We note that even though OGD may be applied to the best expert setting as a special case, it may be shown that $\rho_1(a) = 0$ for $\mathcal{K} = \Delta_N$ or any translation of Δ_N by a constant vector.

(ii) To the L_p -deviation, with $Q_T = \sum_{t=1}^T \|\mathbf{l}_t\|_p^2$ and $\beta = 4$.

For OGD in particular we may apply Theorem 10 in combination with the bound of Hazan and Kale (2010) and the results of Sect. 5.1.⁵ We obtain the following:

Corollary 3 Let $B(\mathbf{0}, a) \subseteq \mathcal{K} \subseteq B(\mathbf{0}, 1)$ for some $a > 0$. Then taking OGD as algorithm A , it holds that

(i) If $\text{Var}_T \geq V$, then for every $\eta > 0$, $\max_t \{R_{A,t}\} \geq \min\{\frac{a^2}{2\eta}, \frac{1}{2}\eta V\}$, and for $\eta = \frac{a}{\sqrt{V}}$,

$$\max_t \{R_{A,t}\} \geq \frac{a}{2} \cdot \sqrt{V}.$$

(ii) If $\sqrt{V'} \geq 12$, then for any sequence with $\text{Var}_T \leq V'$, it holds that

$$\max_t \{R_{A,t}\} \geq \frac{a^2}{60} \cdot \frac{\text{Var}_T}{\sqrt{V'}}.$$

In particular, if $\text{Var}_T = \Theta(V')$, then $\max_t \{R_{A,t}\} = \Omega(\sqrt{\text{Var}_T})$.

7 Extensions and applications

We have shown lower bounds on the anytime regret of concave potential algorithms. We have further proved a trade-off between upper bounds on the regret of such algorithms and lower bounds on their anytime regret. These results were shown in the setting of online linear optimization with full information, which includes the best expert setting. In this work we did not consider the more general setting of online convex optimization or the bandit setting. Whether or not results in a similar spirit may be obtained for those settings is an interesting open question.

7.1 The scope of our results

Our results apply to algorithms whose weight vector is defined as the gradient of a concave function of the cumulative loss. This mechanism is further modulated by using a fixed learning rate. As has already been argued, this algorithmic family contains the important and widely used class of RFTL algorithms. Nevertheless, it is clear that some important algorithms lie outside it. As explained in the previous section, this holds in particular for the Variation MW algorithm of Hazan and Kale (2010) and at least for some instances of the algorithm of Chiang et al. (2012). It remains an interesting open question whether the bound obtained by the Variation MW algorithm may be achieved by means of concave potential algorithms. The same is true for the bounds given in Chiang et al. (2012). We note, however, that those authors show that if the decision set $\mathcal{K}$ contains a ball centered at the origin, then the regret bound of the OGD algorithm in particular is inferior to the regret bound of their algorithm (Chiang et al. 2012, Lemma 9). Specifically, for any learning rate the regret bound of OGD is $\Omega(\min\{D_2, \sqrt{T}\})$, compared with the $O(\sqrt{D_2})$ regret bound of their algorithm.

⁵ Note a very slight change from the original theorem in the second part. This change is needed since the regret upper bound of OGD holds only if $\sqrt{V'} \geq 12$. We also point out that the second part of Theorem 10 still holds if the regret guarantee of the algorithm requires bounded single-period losses.

7.1.1 The use of fixed learning rates

Learning rates are employed as a means of parameterizing regret minimization algorithms in terms of some quantity of interest, such as the variation of the loss sequence. When an upper bound on that quantity is known in advance, the learning rate is commonly fixed as a function of that upper bound. Thus, in such a scenario, our assumption of a fixed learning rate seems like a minor restriction.

When no upper bound is given, variable learning rates are normally employed, with the so-called “doubling trick” being a special case. It is natural to consider adjusting the analysis leading to Lemma 1 to a setting where $\mathbf{x}_t = \nabla \Phi_{\eta_t}(\mathbf{L}_{t-1})$, for some non-increasing learning rates $\eta_1, \dots, \eta_T$. This may be done using the Taylor expansion of Φ as a function of both the learning rate and the cumulative loss. As a result, the regret, which formerly comprised the non-negative first and second order regret terms (see Theorem 2), now contains an additional non-positive term that depends on the changes in the learning rate. (This term disappears if the learning rate is fixed.) It is not clear how to bound the negative contribution of this new term in a way that would yield an interesting modification of the previous analysis.

7.2 Direct and indirect applications

We conclude with two examples that illustrate the applicability of our results.

The first, which was already mentioned in the introduction, concerns lower bounds on the price of financial instruments. Specifically, our results are used to derive a lower bound on the price of “at the money” European call options in an arbitrage-free market. The lower bound obtained for the price of such options is $\exp(0.1\sqrt{Q}) - 1 \approx 0.1\sqrt{Q}$ for $Q < 0.5$, where Q is the assumed quadratic variation of the stock’s log price ratios. This bound has the same asymptotic behavior as the Black-Scholes price and an upper bound on the price given by DeMarzo et al. (2006).

As the core element of this derivation, one requires an algorithm for the best expert setting whose regret is lower bounded in terms of the relative quadratic variation. By our results, the RWM algorithm may guarantee such a lower bound on the regret at some time t (see Sect. 5.2). A further modification to RWM moves all the weight to the best expert at time t , serving to “lock in” the regret until the final round. As a result, the lower bound on the anytime regret of RWM also holds for the final regret of the modified algorithm. For more details, see Gofer and Mansour (2012) and Gofer (2014b).

A second application is in the study of online linear optimization with *switching costs*, namely, where additional costs (and regret) are incurred for switching between weight vectors. It was shown by Gofer (2014a) that for switching costs that are lower bounded by a norm of the difference between weights, regret cannot be upper bounded given only some upper bound on the quadratic variation of the loss sequence, $\sum_t \|\mathbf{l}_t\|_2^2$. The same claim also holds for relative quadratic variation and for the notions of variation considered by Hazan and Kale (2010) and Chiang et al. (2012). In particular, a non-oblivious adversary that knows the expectation of a learner’s next action may force unbounded expected regret. The same is true for an oblivious adversary vs. a deterministic learner. This result holds for general classes of the online linear optimization setting, including the best expert setting and any decision set that contains a ball around the origin.

For proving the above result, an instance of a concave potential algorithm is employed by the *adversary* to adaptively construct a loss sequence that is difficult for a given learning algorithm A . Our results for concave potential algorithms (and in particular RWM and OGD) are used in showing that A incurs either arbitrarily high switching costs and non-negative

regret, or arbitrarily high regret along with the non-negative switching costs. This application demonstrates that regret lower bounds established for concave potential algorithms may facilitate the derivation of regret lower bounds that hold for *every* algorithm.

Acknowledgments This research was supported in part by the Google Inter-university center for Electronic Markets and Auctions, by a grant from the Israel Science Foundation, by a grant from United States-Israel Binational Science Foundation (BSF), by a grant from the Israeli Ministry of Science (MoS), and by The Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11). This work is partly based on Ph.D. thesis research carried out by the first author at Tel Aviv University.

Appendix 1: Additional claims

Lemma 10 *If $\mathcal{K} \subset \mathbb{R}^N$ is compact, then $h(\mathbf{L}) = \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \mathbf{L}\}$ is continuous on $\mathbb{R}^N$.*

Proof First of all, $h(\mathbf{L})$ is well-defined, because for every $\mathbf{L} \in \mathbb{R}^N$, $\mathbf{u} \cdot \mathbf{L}$ is continuous and $\mathcal{K}$ is compact. Similarly, $\|\cdot\|_2$ is continuous on $\mathcal{K}$, so $D = \max_{\mathbf{u} \in \mathcal{K}} \{\|\mathbf{u}\|_2\}$ is well-defined. Let $\mathbf{L}, \mathbf{L}' \in \mathbb{R}^N$, and let $\mathbf{u}, \mathbf{u}' \in \mathcal{K}$ satisfy $h(\mathbf{L}) = \mathbf{u} \cdot \mathbf{L}$ and $h(\mathbf{L}') = \mathbf{u}' \cdot \mathbf{L}'$. Therefore,

$$\begin{aligned} h(\mathbf{L}') - h(\mathbf{L}) &= \mathbf{u}' \cdot \mathbf{L}' - \mathbf{u} \cdot \mathbf{L} \leq \mathbf{u} \cdot \mathbf{L}' - \mathbf{u} \cdot \mathbf{L} = \mathbf{u} \cdot (\mathbf{L}' - \mathbf{L}) \\ &\leq \|\mathbf{u}\|_2 \cdot \|\mathbf{L}' - \mathbf{L}\|_2 \leq D \|\mathbf{L}' - \mathbf{L}\|_2, \end{aligned}$$

where the second inequality is by the Cauchy-Schwarz inequality. Reversing the roles of $\mathbf{L}$ and $\mathbf{L}'$, we get $h(\mathbf{L}) - h(\mathbf{L}') \leq D \|\mathbf{L} - \mathbf{L}'\|_2$, and together we get $|h(\mathbf{L}') - h(\mathbf{L})| \leq D \|\mathbf{L} - \mathbf{L}'\|_2$, which means that h is continuous. $\square$

Lemma 11 *It holds that $\delta : \mathbb{R}^N \rightarrow \mathbb{R}$, $\delta(\mathbf{v}) = \max_i \{v_i\} - \min_i \{v_i\}$, is a seminorm.*

Proof Positive homogeneity: if $a \geq 0$, then $\delta(a\mathbf{v}) = \max_i \{av_i\} - \min_i \{av_i\} = a \max_i \{v_i\} - a \min_i \{v_i\} = |a| \delta(\mathbf{v})$. Otherwise, $\delta(a\mathbf{v}) = \max_i \{av_i\} - \min_i \{av_i\} = -\min_i \{-av_i\} + \max_i \{-av_i\} = \delta(-a\mathbf{v})$, and as already seen, $\delta(-a\mathbf{v}) = |a| \delta(\mathbf{v})$.

Triangle inequality: given $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$, it holds for every j that $\min_i \{u_i\} + \min_i \{v_i\} \leq u_j + v_j \leq \max_i \{u_i\} + \max_i \{v_i\}$. Thus, $\max_i \{u_i + v_i\} \leq \max_i \{u_i\} + \max_i \{v_i\}$, and $\min_i \{u_i\} + \min_i \{v_i\} \leq \min_i \{u_i + v_i\}$. Together we get

$$\begin{aligned} \delta(\mathbf{u} + \mathbf{v}) &= \max_i \{u_i + v_i\} - \min_i \{u_i + v_i\} \\ &\leq \max_i \{u_i\} + \max_i \{v_i\} - \min_i \{u_i\} - \min_i \{v_i\} \\ &= \delta(\mathbf{u}) + \delta(\mathbf{v}). \end{aligned}$$

$\square$

Appendix 2: Missing proofs

Proof of Theorem 3 We start by observing that

$$\begin{aligned} (-1/\eta) \mathcal{R}^*(-\eta \mathbf{L}) &= (-1/\eta) \sup_{\mathbf{x} \in \mathcal{K}} \{\mathbf{x} \cdot (-\eta \mathbf{L}) - \mathcal{R}(\mathbf{x})\} \\ &= \inf_{\mathbf{x} \in \mathcal{K}} \{\mathbf{x} \cdot \mathbf{L} + \mathcal{R}(\mathbf{x})/\eta\} \\ &= \min_{\mathbf{x} \in \mathcal{K}} \{\mathbf{x} \cdot \mathbf{L} + \mathcal{R}(\mathbf{x})/\eta\}. \end{aligned}$$

We used the fact that $\mathbf{x} \cdot \mathbf{L} + \mathcal{R}(\mathbf{x})/\eta$ is continuous and attains a minimum on a compact set. Consequently, Φ is defined for every $\mathbf{L}$. Since $\mathcal{R}^*$ is convex, $\Phi(\mathbf{L}) = (-1/\eta)\mathcal{R}^*(-\eta\mathbf{L})$ is concave.

Denote $h(\mathbf{u}, \mathbf{L}) = \mathbf{u} \cdot \mathbf{L} + \mathcal{R}(\mathbf{u})/\eta$, and let $\mathbf{L}, \mathbf{L}' \in \mathbb{R}^N$, where $\mathbf{L} \neq \mathbf{L}'$. In addition, let $\mathbf{x} = \arg \min_{\mathbf{u} \in \mathcal{K}} \{h(\mathbf{u}, \mathbf{L})\}$ and $\mathbf{x}' = \arg \min_{\mathbf{u} \in \mathcal{K}} \{h(\mathbf{u}, \mathbf{L}')\}$, and denote $\Delta\mathbf{L} = \mathbf{L}' - \mathbf{L}$ and $\Delta\mathbf{x} = \mathbf{x}' - \mathbf{x}$. We have

$$\Phi(\mathbf{L}') - \Phi(\mathbf{L}) = h(\mathbf{x}', \mathbf{L}') - h(\mathbf{x}, \mathbf{L}) \leq h(\mathbf{x}, \mathbf{L}') - h(\mathbf{x}, \mathbf{L}) = \mathbf{x} \cdot \Delta\mathbf{L},$$

where the inequality is by definition of $\mathbf{x}'$. Reversing the roles of $\mathbf{L}$ and $\mathbf{L}'$, we also get $\Phi(\mathbf{L}) - \Phi(\mathbf{L}') \leq -\mathbf{x}' \cdot \Delta\mathbf{L}$. Therefore, $\mathbf{x}' \cdot \Delta\mathbf{L} \leq \Phi(\mathbf{L}') - \Phi(\mathbf{L}) \leq \mathbf{x} \cdot \Delta\mathbf{L}$, or $\Delta\mathbf{x} \cdot \Delta\mathbf{L} \leq \Phi(\mathbf{L}') - \Phi(\mathbf{L}) - \mathbf{x} \cdot \Delta\mathbf{L} \leq 0$, implying that

$$|\Phi(\mathbf{L}') - \Phi(\mathbf{L}) - \mathbf{x} \cdot \Delta\mathbf{L}| \leq |\Delta\mathbf{x} \cdot \Delta\mathbf{L}| \leq \|\Delta\mathbf{x}\|_2 \|\Delta\mathbf{L}\|_2,$$

or $\frac{|\Phi(\mathbf{L} + \Delta\mathbf{L}) - \Phi(\mathbf{L}) - \mathbf{x} \cdot \Delta\mathbf{L}|}{\|\Delta\mathbf{L}\|_2} \leq \|\Delta\mathbf{x}\|_2$. If we showed that $\lim_{\Delta\mathbf{L} \rightarrow \mathbf{0}} \|\Delta\mathbf{x}\|_2 = 0$, then

$$\lim_{\Delta\mathbf{L} \rightarrow \mathbf{0}} \frac{\Phi(\mathbf{L} + \Delta\mathbf{L}) - \Phi(\mathbf{L}) - \mathbf{x} \cdot \Delta\mathbf{L}}{\|\Delta\mathbf{L}\|_2} = 0,$$

that is, $\nabla\Phi(\mathbf{L}) = \mathbf{x}$, as required. Furthermore, the fact that Φ is differentiable on $\mathbb{R}^N$ would imply that it is continuously differentiable (see, e.g., Corollary 25.5.1 in [Rockafellar 1970](#)).

Since $\mathcal{R}$ is strongly convex (with some parameter $\alpha > 0$), we have that $\mathcal{R}(\frac{1}{2}\mathbf{x} + \frac{1}{2}\mathbf{x}') \leq \frac{1}{2}\mathcal{R}(\mathbf{x}) + \frac{1}{2}\mathcal{R}(\mathbf{x}') - \frac{\alpha}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} \|\Delta\mathbf{x}\|_2^2$. Dividing by η and adding $\mathbf{L} \cdot (\frac{1}{2}\mathbf{x} + \frac{1}{2}\mathbf{x}')$ to both sides, we get

$$\begin{aligned} \frac{1}{\eta} \mathcal{R}\left(\frac{\mathbf{x} + \mathbf{x}'}{2}\right) + \mathbf{L} \cdot \left(\frac{\mathbf{x} + \mathbf{x}'}{2}\right) &\leq \frac{1}{2\eta} \mathcal{R}(\mathbf{x}) + \mathbf{L} \cdot \frac{1}{2}\mathbf{x} + \frac{1}{2\eta} \mathcal{R}(\mathbf{x}') \\ &\quad + \mathbf{L} \cdot \frac{1}{2}\mathbf{x}' - \frac{\alpha}{8\eta} \|\Delta\mathbf{x}\|_2^2, \end{aligned}$$

or

$$h\left(\frac{\mathbf{x} + \mathbf{x}'}{2}, \mathbf{L}\right) \leq \frac{1}{2}h(\mathbf{x}, \mathbf{L}) + \frac{1}{2}h(\mathbf{x}', \mathbf{L}) - \frac{\alpha}{8\eta} \|\Delta\mathbf{x}\|_2^2.$$

By definition of $\mathbf{x}$, $h\left(\frac{\mathbf{x} + \mathbf{x}'}{2}, \mathbf{L}\right) \geq h(\mathbf{x}, \mathbf{L})$, so, rearranging, we get

$$h(\mathbf{x}', \mathbf{L}) - h(\mathbf{x}, \mathbf{L}) \geq \frac{\alpha}{4\eta} \|\Delta\mathbf{x}\|_2^2.$$

Since

$$\begin{aligned} h(\mathbf{x}', \mathbf{L}) &= \mathbf{x}' \cdot \mathbf{L} + \frac{1}{\eta} \mathcal{R}(\mathbf{x}') = \mathbf{x}' \cdot \mathbf{L}' + \frac{1}{\eta} \mathcal{R}(\mathbf{x}') - \mathbf{x}' \cdot \Delta\mathbf{L} \\ &= h(\mathbf{x}', \mathbf{L}') - \mathbf{x}' \cdot \Delta\mathbf{L} \leq h(\mathbf{x}, \mathbf{L}') - \mathbf{x}' \cdot \Delta\mathbf{L}, \end{aligned}$$

we have that

$$\begin{aligned} \frac{\alpha}{4\eta} \|\Delta\mathbf{x}\|_2^2 &\leq h(\mathbf{x}, \mathbf{L}') - \mathbf{x}' \cdot \Delta\mathbf{L} - h(\mathbf{x}, \mathbf{L}) \\ &= \mathbf{x} \cdot \mathbf{L}' - \mathbf{x}' \cdot \Delta\mathbf{L} - \mathbf{x} \cdot \mathbf{L} = -\Delta\mathbf{x} \cdot \Delta\mathbf{L} \\ &\leq \|\Delta\mathbf{x}\|_2 \|\Delta\mathbf{L}\|_2, \end{aligned}$$

and thus, $\|\Delta\mathbf{x}\|_2 \leq \frac{4\eta}{\alpha} \|\Delta\mathbf{L}\|_2$. Therefore, $\lim_{\Delta\mathbf{L} \rightarrow \mathbf{0}} \|\Delta\mathbf{x}\|_2 = 0$, completing the proof.

Proof of Lemma 6 We assume w.l.o.g. that $p_{i,0} > 0$ for every i . Let $1 \leq i, j \leq N$. We have that

$$\frac{\partial}{\partial L_i} \left(\ln \frac{\partial \Phi_\eta}{\partial L_j} \right) = \frac{\partial^2 \Phi_\eta}{\partial L_i \partial L_j} \cdot \left(\frac{\partial \Phi_\eta}{\partial L_j} \right)^{-1} = \frac{\partial^2 \Phi_\eta}{\partial L_i \partial L_j} p_j^{-1}.$$

Since

$$\ln \frac{\partial \Phi_\eta}{\partial L_j} = \ln \left(\frac{p_{i,0} e^{-\eta L_j}}{\sum_{k=1}^N p_{k,0} e^{-\eta L_k}} \right) = \ln p_{i,0} - \eta L_j + \eta \Phi_\eta,$$

we also have that

$$\frac{\partial}{\partial L_i} \left(\ln \frac{\partial \Phi_\eta}{\partial L_j} \right) = -\eta \frac{\partial L_j}{\partial L_i} + \eta \frac{\partial \Phi_\eta}{\partial L_i} = -\eta \cdot 1_{i=j} + \eta p_i,$$

and therefore, $\frac{\partial^2 \Phi_\eta}{\partial L_i \partial L_j} p_j^{-1} = -\eta \cdot 1_{i=j} + \eta p_i$, or, equivalently, $\frac{\partial^2 \Phi_\eta}{\partial L_i \partial L_j} = \eta(p_i p_j - 1_{i=j} p_j)$, as required.

For any $\mathbf{p} \in \Delta_N$ and $\mathbf{z} \in \mathbb{R}^N$, we denote by $X_{\mathbf{p}, \mathbf{z}}$ a random variable that satisfies $\mathbb{P}(X_{\mathbf{p}, \mathbf{z}} = z_i) = \sum_{\{j: z_j = z_i\}} p_j$ for every $1 \leq i \leq N$. Note that even though the values z_i are not necessarily distinct, nevertheless, $\mathbb{E}[X_{\mathbf{p}, \mathbf{z}}] = \sum_{i=1}^N p_i z_i$ and $\text{Var}(X_{\mathbf{p}, \mathbf{z}}) = \sum_{i=1}^N p_i z_i^2 - (\sum_{i=1}^N p_i z_i)^2$ as usual, by continuity. Thus, for any $\mathbf{z} \in \mathbb{R}^N$, we have that

$$\begin{aligned} 0 \leq \text{Var}(X_{\mathbf{p}, \mathbf{z}}) &= \sum_{i=1}^N p_i z_i^2 - \left(\sum_{i=1}^N p_i z_i \right)^2 = \mathbf{z}^\top \text{diag}(\mathbf{p}) \mathbf{z} - \mathbf{z}^\top \mathbf{p} \mathbf{p}^\top \mathbf{z} \\ &= \mathbf{z}^\top (\text{diag}(\mathbf{p}) - \mathbf{p} \mathbf{p}^\top) \mathbf{z}. \end{aligned}$$

We therefore have that $\text{diag}(\mathbf{p}) - \mathbf{p} \mathbf{p}^\top = -(1/\eta) \nabla^2 \Phi_\eta(\mathbf{L})$ is positive semi-definite, and the result follows.

Proof of Lemma 8 Denoting $y_k = x_k - \mathbb{E}[X]$ for $1 \leq k \leq N$ and $p = \frac{p_i}{p_i + p_j}$, we have that

$$\begin{aligned} \text{Var}(X) &\geq p_i y_i^2 + p_j y_j^2 \geq p_i y_i^2 + p_j y_j^2 - (p_i + p_j)(p y_i + (1-p) y_j)^2 \\ &= (p_i + p_j)(p y_i^2 + (1-p) y_j^2 - (p y_i + (1-p) y_j)^2) \\ &= (p_i + p_j)((p - p^2) y_i^2 + (p - p^2) y_j^2 - 2(p - p^2) y_i y_j) \\ &= (p_i + p_j)(p - p^2)(y_i - y_j)^2 = \frac{p_i p_j}{p_i + p_j} (y_i - y_j)^2 \\ &= \frac{p_i p_j}{p_i + p_j} (x_i - x_j)^2. \end{aligned}$$

To obtain an equality, the first inequality must become an equality, necessitating that $x_k = \mathbb{E}[X]$ for every $k \neq i, j$, and thus, $N \leq 3$. If $N = 2$, the first inequality is an equality, and $p y_i + (1-p) y_j = p_1 x_1 + p_2 x_2 - \mathbb{E}[X] = 0$, so the second above inequality is also an equality. For $N = 3$, if $x_k = \mathbb{E}[X]$, then the first inequality becomes an equality, and in addition,

$$\begin{aligned} p y_i + (1-p) y_j &= \frac{p_i x_i + p_j x_j}{p_i + p_j} - \mathbb{E}[X] = \frac{\mathbb{E}[X] - p_k x_k}{p_i + p_j} - \mathbb{E}[X] \\ &= \frac{\mathbb{E}[X](1 - p_k)}{p_i + p_j} - \mathbb{E}[X] = 0, \end{aligned}$$

so the second inequality becomes an equality. Thus, we have an equality iff $x_k = \mathbb{E}[X]$, or equivalently, $i = 1, k = 2, j = 3$, and $x_2 = p_1x_1 + p_2x_2 + p_3x_3$.

Finally, we have in particular that

$$\begin{aligned} \text{Var}(X) &\geq \frac{p_1 p_N}{p_1 + p_N} \cdot (x_N - x_1)^2 = \frac{1}{1/p_1 + 1/p_N} \cdot (x_N - x_1)^2 \\ &\geq \frac{1}{2} \min_k \{p_k\} (x_N - x_1)^2, \end{aligned}$$

completing the proof.

Appendix 3: Extending the FTL–BTL lemma to lower bound regret

The upper bound on the regret of the RFTL algorithm relies on the “follow the leader, be the leader” (FTL–BTL) lemma (see Hazan 2011). We extend that lemma to facilitate the proof of a lower bound as well. The following lemma gives the core inductive argument of the extension as well as the original, for completeness.

Lemma 12 *Let X be a non-empty set, let $1 \leq n \in \mathbb{N}$, and let $g_i : X \rightarrow \mathbb{R}$ for $1 \leq i \leq n$. Denote $G_i = \sum_{j=1}^i g_j$, and assume G_i has at least one global minimum on X . Let x_i be such a minimum for $1 \leq i \leq n$, and let $x_0 \in X$ be an arbitrary point. Then $\sum_{j=1}^n g_j(x_j) \leq G_n(x_n) \leq \sum_{j=1}^n g_j(x_{j-1})$.*

Proof By induction on n . For $n = 1$, $\sum_{j=1}^n g_j(x_j) = G_1(x_1) \leq G_1(x_0) = g_1(x_0) = \sum_{j=1}^n g_j(x_{j-1})$, where the inequality is by definition of x_1 . We assume correctness for n and proceed to prove for $n + 1$. We have on the one hand that

$$\begin{aligned} G_{n+1}(x_{n+1}) &\leq G_{n+1}(x_n) = G_n(x_n) + g_{n+1}(x_n) \\ &\leq \sum_{j=1}^n g_j(x_{j-1}) + g_{n+1}(x_n) = \sum_{j=1}^{n+1} g_j(x_{j-1}), \end{aligned}$$

where the first inequality is by definition of x_{n+1} , and the second one is by the induction assumption. On the other hand,

$$\begin{aligned} G_{n+1}(x_{n+1}) &= G_n(x_{n+1}) + g_{n+1}(x_{n+1}) \geq G_n(x_n) + g_{n+1}(x_{n+1}) \\ &\geq \sum_{j=1}^n g_j(x_j) + g_{n+1}(x_{n+1}) = \sum_{j=1}^{n+1} g_j(x_j), \end{aligned}$$

where the first inequality is by definition of x_n , and the second inequality is by the induction assumption. $\square$

We now give the extended FTL–BTL lemma. We assume the RFTL setting, so that $\mathcal{K} \subset \mathbb{R}^N$ is non-empty, compact, and convex, $\mathcal{R} : \mathcal{K} \rightarrow \mathbb{R}$ is strongly convex, and $\eta > 0$.

Lemma 13 *Let $\mathbf{l}_t \in \mathbb{R}^N$ for $1 \leq t \leq T$, and let $\mathbf{u}_0 = \arg \min_{\mathbf{u} \in \mathcal{K}} \{\mathbf{u} \cdot \sum_{t=1}^T \mathbf{l}_t\}$. Then*

$$R_{\text{RFTL}(\eta, \mathcal{R}), T} \leq \sum_{t=1}^T \mathbf{l}_t \cdot (\mathbf{x}_t - \mathbf{x}_{t+1}) + (\mathcal{R}(\mathbf{u}_0) - \mathcal{R}(\mathbf{x}_1))/\eta$$

and

$$R_{RFTL(\eta, \mathcal{R}), T} \geq (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1))/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot (\mathbf{x}_{T+1} - \mathbf{u}_0) \geq 0.$$

Proof To fit the terminology of Lemma 12, set $X = \mathcal{K}$, $n = T + 1$, $g_1(\mathbf{x}) = \mathcal{R}(\mathbf{x})/\eta$, and $g_{t+1}(\mathbf{x}) = \mathbf{l}_t \cdot \mathbf{x}$ for $1 \leq t \leq T$. Thus, $G_{t+1}(\mathbf{x}) = \mathcal{R}(\mathbf{x})/\eta + \mathbf{x} \cdot \sum_{\tau=1}^t \mathbf{l}_\tau$ for $0 \leq t \leq T$, and $\mathbf{x}_{t+1} = \arg \min_{\mathbf{u} \in \mathcal{K}} \{\mathcal{R}(\mathbf{u})/\eta + \mathbf{u} \cdot \sum_{\tau=1}^t \mathbf{l}_\tau\}$. In addition, we set $\mathbf{x}_0 = \mathbf{x}_1$. By Lemma 12,

$$\mathcal{R}(\mathbf{x}_1)/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{t+1} \leq \mathcal{R}(\mathbf{x}_{T+1})/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{T+1} \leq \mathcal{R}(\mathbf{x}_1)/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_t,$$

and thus

$$\sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{t+1} - \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_t \leq (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1))/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{T+1} - \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_t \leq 0.$$

We focus first on the left inequality. By definition of $\mathbf{x}_{T+1}$, $G_{T+1}(\mathbf{x}_{T+1}) \leq G_{T+1}(\mathbf{u}_0)$, implying that $\mathcal{R}(\mathbf{x}_{T+1})/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{T+1} \leq \mathcal{R}(\mathbf{u}_0)/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{u}_0$. Therefore, $\sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{t+1} - \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_t \leq (\mathcal{R}(\mathbf{u}_0) - \mathcal{R}(\mathbf{x}_1))/\eta + \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{u}_0 - \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_t$, and rearranging,

$$R_{RFTL(\eta, \mathcal{R}), T} \leq \sum_{t=1}^T \mathbf{l}_t \cdot (\mathbf{x}_t - \mathbf{x}_{t+1}) + (\mathcal{R}(\mathbf{u}_0) - \mathcal{R}(\mathbf{x}_1))/\eta,$$

which is the known upper bound expression on the regret of RFTL.

Moving on to the other inequality, we have that

$$\begin{aligned} R_{RFTL(\eta, \mathcal{R}), T} &= \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_t - \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{u}_0 \\ &\geq (\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1))/\eta + \left(\sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{T+1} - \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{u}_0 \right). \end{aligned}$$

It holds that $(\mathcal{R}(\mathbf{x}_{T+1}) - \mathcal{R}(\mathbf{x}_1))/\eta \geq 0$ by definition of $\mathbf{x}_1$, and $\sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{x}_{T+1} - \sum_{t=1}^T \mathbf{l}_t \cdot \mathbf{u}_0 \geq 0$ by definition of $\mathbf{u}_0$. $\square$

References

- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge: Cambridge University Press.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge: Cambridge University Press.
- Cesa-Bianchi, N., Mansour, Y., & Stoltz, G. (2007). Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2–3), 321–352.
- Chiang, C. K., Yang, T., Lee, C. J., Mahdavi, M., Lu, C. J., Jin, R., et al. (2012). Online optimization with gradual variations. *Journal of Machine Learning Research: Proceedings Track*, 23, 6.1–6.20.
- DeMarzo, P., Kremer, I., & Mansour, Y. (2006). Online trading algorithms and robust option pricing. In *Proceedings of the 38th annual ACM symposium on theory of computing* (pp. 477–486).
- Even-Dar, E., Kearns, M., Mansour, Y., & Wortman, J. (2008). Regret to the best vs. regret to the average. *Machine Learning*, 72(1–2), 21–37.
- Gofer, E. (2014a). Higher-order regret bounds with switching costs. *Journal of Machine Learning Research: Proceedings Track*, 35, 210–243.
- Gofer, E. (2014b). *Machine learning algorithms with applications in finance*. PhD thesis, Tel Aviv University.

- Gofer, E., & Mansour, Y. (2012). Lower bounds on individual sequence regret. In N. H. Bshouty, G. Stoltz, N. Vayatis, & T. Zeugmann (Eds.), *Algorithmic learning theory* (pp. 275–289). Heidelberg: Springer.
- Hazan, E. (2006). *Efficient algorithms for online convex optimization and their applications*. PhD thesis, Princeton University.
- Hazan, E. (2011). The convex optimization approach to regret minimization. In S. Sra, S. Nowozin, & S. J. Wright (Eds.), *Optimization for machine learning*. Cambridge: MIT Press.
- Hazan, E., & Kale, S. (2010). Extracting certainty from uncertainty: Regret bounded by variation in costs. *Machine Learning*, 80(2–3), 165–188.
- Kalai, A., & Vempala, S. (2005). Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3), 291–307.
- Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*. Dordrecht: Kluwer Academic Publishers.
- Rockafellar, R. T. (1970). *Convex analysis*. Princeton: Princeton University Press.
- Shalev-Shwartz, S. (2007). *Online learning: Theory, algorithms, and applications*. PhD thesis, The Hebrew University.
- von Szökefalvi Nagy, J. (1918). Über algebraische gleichungen mit lauter reellen wurzeln. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 27, 37–43.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *ICML* (pp. 928–936).