



HAL
open science

Efficient and Near-Optimal Online Portfolio Selection

Rémi Jézéquel, Dmitrii M. Ostrovskii, Pierre Gaillard

► **To cite this version:**

Rémi Jézéquel, Dmitrii M. Ostrovskii, Pierre Gaillard. Efficient and Near-Optimal Online Portfolio Selection. 2022. hal-03787674

HAL Id: hal-03787674

<https://hal.science/hal-03787674v1>

Preprint submitted on 27 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficient and Near-Optimal Online Portfolio Selection

Rémi Jézéquel^{*†}

Dmitrii M. Ostrovskii^{*‡}

Pierre Gaillard[§]

Abstract

In the problem of online portfolio selection as formulated by Cover (1991) [12], the trader repeatedly distributes her capital over d assets in each of $T > 1$ rounds, with the goal of maximizing the total return. Cover proposed an algorithm, termed Universal Portfolios, that performs nearly as well as the best (in hindsight) static assignment of a portfolio, with an $O(d \log(T))$ logarithmic regret. Without imposing any restrictions on the market this guarantee is known to be worst-case optimal, and no other algorithm attaining it has been discovered so far. Unfortunately, Cover’s algorithm crucially relies on computing certain d -dimensional integral, which must be approximated in any implementation; this results in a prohibitive $\tilde{O}(d^4(T+d)^{14})$ per-round runtime for the fastest known implementation due to Kalai and Vempala (2002) [22]. We propose an algorithm for online portfolio selection that admits essentially the same regret guarantee as Universal Portfolios—up to a constant factor and replacement of $\log(T)$ with $\log(T+d)$ —yet has a drastically reduced runtime of $\tilde{O}(d^2(T+d))$ per round. The selected portfolio minimizes the observed logarithmic loss regularized with the log-determinant of its Hessian—equivalently, the hybrid logarithmic-volumetric barrier of the polytope specified by the asset return vectors. As such, our work reveals surprising connections of online portfolio selection with two classical topics in optimization theory: cutting-plane and interior-point algorithms.

1 Introduction

In the problem of portfolio selection in its most general form, the goal is to allocate capital over a set of assets in such a way as to maximize the aggregate return; see e.g. [32, 43] and references therein. At a high level, one can distinguish two rather different approaches to this general problem.

- (a) Markowitz’s *mean-variance* theory [31] in which the tradeoff between the expected return and its variance allows to account for the market uncertainty in a one-time investment scenario.
- (b) The *capital growth theory* ascending to [23]—see [30] for a modern treatment—which focuses on the expected *log return* of a portfolio and naturally addresses multiple-period investment.

We shall focus here on the *online portfolio selection* framework introduced by T. Cover [12, 14]. While being close in spirit to the second approach and addressing the case of multi-period investment, Cover’s framework allows to avoid making any stochastic assumptions about the market.

^{*}Equal contribution.

[†]Inria & Département d’Informatique de l’École Normale Supérieure, Paris, France. remi.jezequel@inria.fr.

[‡]Department of Mathematics, University of Southern California, Los Angeles, USA. dostrovs@usc.edu.

[§]Inria Grenoble-Rhône-Alpes & Université Grenoble Alpes, Grenoble, France. pierre.gaillard@inria.fr.

This results, on the one hand, in a robust theory well suited for dynamic environment and, on the other hand, in a plethora of algorithms with great practical performance and strong guarantees [28].

Cover’s framework can be summarized as follows. Given the unit initial capital $\text{Cap}_0 = 1$, in each round $t \in \{1, 2, \dots, T\}$ the trader chooses an allocation $w_t \in \Delta_d$ of her previously earned capital Cap_{t-1} over d assets; here Δ_d is the standard simplex in \mathbb{R}^d and $w_t[i]$ —the i^{th} entry of w_t —is the share of capital invested into asset $i \in \{1, 2, \dots, d\}$ at this round. Then the *returns*—the ratios of the closing and opening prices in this round—are revealed in the form of $x_t \in \mathbb{R}_+^d$ (hereafter \mathbb{R}_+^d is the nonnegative orthant of \mathbb{R}^d , and \mathbb{R}_{++}^d is its interior), and the trader’s capital is updated as

$$\text{Cap}_t = \text{Cap}_{t-1} x_t^\top w_t.$$

By Cover, the performance of a strategy that selected portfolios $w_{1:T} := (w_1, \dots, w_T)$ for the market given by $x_{1:T} := (x_1, \dots, x_T)$ is quantified by comparing the final capital $\text{Cap}_T = \prod_{t=1}^T x_t^\top w_t$ against

$$\text{Cap}_T^o := \max_{w \in \Delta_d} \prod_{t=1}^T x_t^\top w,$$

the “idealized” final capital attained by the best (in hindsight for $x_{1:T}$) “static” strategy constrained to select the same portfolio in all rounds. Due to the multiplicative structure of Cap_T and Cap_T^o , it is natural to define the *regret* of $w_{1:T}$ on market $x_{1:T}$ as the negative logarithm of $\text{Cap}_T/\text{Cap}_T^o$, i.e.

$$\mathcal{R}_T(w_{1:T}|x_{1:T}) := \sum_{t \in [T]} \ell_t(w_t) - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w) \quad (1)$$

where we let $[T] := \{1, 2, \dots, T\}$ for $T \in \mathbb{N}$ with the convention $[0] = \emptyset$, and define the instantaneous (logarithmic) loss $\ell_t : \Delta_d \rightarrow \mathbb{R}$ by

$$\ell_t(w) := -\log(x_t^\top w). \quad (2)$$

Since the instantaneous losses are convex, this formulation falls into the framework of online convex optimization, with (1) matching the standard notion of regret for an online optimization algorithm [19, Sec. 1.1]; thus, we can benefit from an abundance of techniques developed by the corresponding community [19]. Despite conveniently putting us into the framework of online convex optimization, comparison against the class of “static” strategies, also called *constantly rebalanced portfolios* (CRP), might seem restrictive. Yet, there are at least two other arguments in its favor.

- On the one hand, CRP strategies are known to be optimal in the i.i.d. stochastic setup where x_t ’s, instead of being adversarially chosen in each round, come from a static distribution unknown to the trader; see [13, Theorem 15.3.1]. Clearly, a uniform upper bound for the regret (1), i.e. one valid for any sequence $x_{1:T}$, also applies in the i.i.d. setup. Such models are adequate for practical markets in the case of moderate observation horizons; hence, we may expect good *practical* performance for a strategy admitting a uniform regret bound against the best CRP. This intuition is verified by extensive real-data experiments reported in [27].
- On the other hand, the regret measured against the best *sequence of portfolios* can be as large as $T \log(d)$,¹ i.e. grow linearly with time horizon T . Meanwhile, the regret measured against the best CRP strategy, cf. (1), is expected to be *sublinear in T* for any reasonable algorithm; in fact, this is the case for all algorithms we shall discuss from now on.

¹Indeed: given any $w_t \in \Delta_d$, let $x_t = e_{i(t)}$ with $i(t) \in \text{Argmin}_{1 \leq i \leq d} w_t[i]$ where $\{e_1, \dots, e_d\}$ is the canonical basis of \mathbb{R}^d . Then $\log(\Pi_T) = \log(\prod_{t \leq T} w_t[i(t)]) \leq \log[(\frac{1}{d})^T] = -T \log d$, but at the same time $\ell_t(w)$ vanishes on $e_{i(t)} \in \Delta_d$.

In addition to introducing the framework, in his seminal work [12] Cover proposed and studied an algorithm for online portfolio selection, termed *Universal Portfolios*, which consists in selecting

$$w_t = \int_{\Delta_d} w \phi_t(w) dw \quad (3)$$

where $\phi_t(w)$ is the probability density on Δ_d proportional to the capital the trader would have earned by the current round t by selecting portfolio w in all the previous rounds—in other words,

$$\phi_t(w) \propto \exp \left(- \sum_{\tau \in [t-1]} \ell_\tau(w) \right). \quad (4)$$

The idea here is to give preference to portfolios for which the observed cumulative loss is relatively small without completely discarding portfolios with large losses. In fact, the same procedure under different names (e.g., the method of exponential weights, exponentially weighted averaging forecaster) arises in other contexts: prediction with expert advice [11], online learning in metric spaces [41], and statistical estimation [26, 16, 15, 42, 5, 6]. Cover justified his algorithm by showing that for any market realization $x_{1:T}$, the regret of portfolio sequence $w_{1:T}$ produced in (3) satisfies

$$\mathcal{R}_T(w_{1:T}|x_{1:T}) = (d-1) \log(T+1) \quad (5)$$

see [12, Thm. 4] and [11, Thm. 10.3].² This result justifies the name “Universal Portfolios:” (3) admits a strong—nearly independent of T —regret guarantee for *any* market. Moreover, this regret bound turns out to be optimal up to a constant factor: it can be shown that the worst-case over $x_{1:T}$ regret for *any* algorithm is at least $\frac{d-1}{2} \log(\frac{\pi T}{d}) + \varepsilon_d(T)$ where $\varepsilon_d(T) \rightarrow 0$ as $T \rightarrow \infty$; see [11, p. 282].

Unfortunately, Cover’s algorithm has a major drawback: it relies on integration over Δ_d , which cannot be performed explicitly, and thus has to be approximated with sufficient accuracy to obtain a tractable implementation. A natural approach would be to sample from ϕ_t —a log-concave distribution—using some Markov Chain Monte Carlo techniques. This approach was carried out by Kalai and Vempala in [22], resulting in a $\tilde{O}(T^5(T+d)^9 d^4)$ per-round runtime.³ (Hereinafter $\tilde{O}(\cdot)$ hides a polylogarithmic factor.) Albeit polynomial in d and T , this runtime essentially rules out any practical application scenario. Even though this can be improved via the recent advances in log-concave sampling (see e.g. [9, 35, 17, 10, 48]), the dependency on T is likely to remain prohibitive. Meanwhile, none of the algorithms proposed since Cover’s original work is known to admit a regret bound matching (5) up to a constant factor. In other words, the following problem stands:

*Propose a **computationally feasible** and **near-optimal** online portfolio selection algorithm.*

We shall address this problem by presenting such an algorithm, proving that its regret admits a bound matching (5) up to a constant factor and replacement of $\log(T)$ with $\log(T+d)$, and proposing an efficient implementation for it. In the remaining part of this section, we shall informally stating our main result, briefly summarize our contributions, discuss the related work, and give an outline of the subsequent sections of the paper. Let us now give an abridged formulation of our main result.

²See also [20, Theorem 7] for a short self-contained proof of this upper bound with an additional constant term.

³In [46] this result is miscited as $\tilde{O}(T^3(T+d)d^2)$, seemingly due to neglecting that $\frac{1}{\delta_0^2} = T(T+d)^2 d$ in [22, Thm. 2].

Theorem 1.1 (cf. Theorem 4.1). *There is an algorithm for online portfolio selection which has an $O(d \log(T + d))$ regret, runs in $\tilde{O}(d^2 T + d^3)$ per round, and utilizes $O(dT + d^2)$ memory.*

We see that the improvement of runtime compared to the algorithm of Kalai and Vempala [22]—the only one known so far with a regret guarantee of the same order—is drastic; in particular, our proposed algorithm can be run on a personal computer for d and T in the range of thousands. Moreover, in addition to favorable computational guarantees the algorithm is conceptually easy to implement: in each round one has to minimize a convex potential function with a readily available second-order oracle, so one can simply use Newton’s method. In fact, the potential we use belongs to the class of *self-concordant* functions—“canonical” objectives for Newton’s method [37]; this allows us to implement the algorithm while performing only $O(\log(T))$ Newton steps in each round.

We shall now briefly summarize the contributions of our paper and outline its organization.

Contributions and outline of further sections. In Section 2 we present our key contributions.

- We introduce a new algorithm for online portfolio selection, called VB-FTRL—the acronym stands for **Volumetric-Barrier enhanced Follow-The-Regularized-Leader**,—and state Theorem 2.1 showing that the regret of VB-FTRL admits an $O(d \log(T + d))$ upper bound.
- We then investigate an intriguing question of how VB-FTRL is related to Cover’s Universal Portfolios algorithm. To this end, we first pass to a variational formulation of Cover’s update by finding a functional over distributions supported on Δ_d which is minimized by $\phi_t(\cdot)$, cf. (4). We then show that the update of VB-FTRL corresponds to solving this minimization problem approximately, by passing to the class of appropriately truncated Gaussian distributions with unknown expectation and covariance, and solving for covariance with a given expectation. Moreover, we quantify the accuracy of this approximation (Propositions 2.2 and 2.3); this result might be of independent interest and find use beyond the context of portfolio selection.

Section 3 is devoted to analyzing the regret of VB-FTRL and proving Theorem 2.1. Finally, in Section 4 we present a procedure implementing VB-FTRL, and quantify its time and memory costs.

1.1 Overview of related work

Since Cover’s work [12], the online portfolio selection framework has been studied extensively due to its practical usefulness and theoretical challenges. Here we focus on the line of research done in the online convex optimization community and concerned with designing efficient methods with provable regret guarantees; for reader’s convenience we recap these results in Table 1. For a more application-centered perspective and experimental comparison of algorithms one may refer to [28].

To the best of our knowledge, the authors of [21] were the first to realize that Cover’s problem of competing with the best CRP can be reformulated as online convex optimization on the simplex Δ_d with logarithmic instantaneous loss $\ell_t(w) = -\log(x_t^\top w)$, cf. (1). They advocated the exponentiated gradient (EG) method [25] over the basic online gradient descent [50], due to EG being better adapted to the simplex type geometry. Unfortunately, both methods require the loss gradients to

Algorithm	Regret	Runtime (per round)	Sources
Universal Portfolios	$d \log(T)$	$d^4 T^{14}$	[12, 22]
Online Gradient Descent	$G_2 \sqrt{T}$	d	[50]
Exponentiated Gradient	$G_\infty \sqrt{T \log(d)}$	d	[25, 21]
Online Newton Step (ONS)	$G_\infty d \log(T)$	$d^2 +$ generalized projection on Δ_d^4	[1, 20]
Soft-Bayes	$\sqrt{dT \log(d)}$	d	[39]
Ada-BARRONS	$d^2 \log^4(T)$	$d^{2.5} T$	[29]
BISONS	$d^2 \log^2(T)$	$\text{poly}(d)$	[49]
AdaMix+DONS	$d^2 \log^5(T)$	d^3	[34]
VB-FTRL	$d \log(T)$	$d^2 T$	our paper

Table 1: Regret guarantees and per-round runtime for various online portfolio selection algorithms. G_p is an upper bound on $\max_{1 \leq t \leq T} \|\nabla \ell_t(w)\|_p$ whenever such a bound is assumed. For brevity we omit a constant factor for the regret and a polylogarithmic factor for the runtime, and assume $T \geq d$.

be bounded, which is a restrictive assumption. Indeed, letting $\|\cdot\|_p$ be ℓ_p -norm on \mathbb{R}^d we have that

$$\|\nabla \ell_t(w)\|_p = \frac{\|x_t\|_p}{x_t^\top w}. \quad (6)$$

For EG, the assumption is that $\|\nabla \ell_t(w)\|_\infty \leq G_\infty$ for some $G_\infty < \infty$; by (6) this can be enforced in either of two ways: (a) by restricting the set of portfolios to $\{w \in \Delta_d : \min_i w[i] \geq 1/G_\infty\}$, or (b) by bounding the dynamic range of asset returns, i.e. requiring that $\frac{\max_i x_t[i]}{\min_i x_t[i]} \leq G_\infty$ for all $1 \leq t \leq T$.

Another issue with these methods is that the regret is suboptimal in T , growing as \sqrt{T} , cf. Table 1.

The two issues afflicting the previous methods were addressed—in isolation—in [39] and [1, 20]. Namely, the Soft-Bayes algorithm from [39] achieves an $O(\sqrt{dT \log(d)})$ regret without assuming boundedness of the loss gradients. Meanwhile, Online Newton Step (ONS), an algorithm proposed in [1, 20], has an $O(G_\infty d \log(T))$ regret under the same assumption of ℓ_∞ -bounded loss gradients as for EG, which matches the guarantee (5) for Universal Portfolios in the regime of a constant G_∞ . All these algorithms can be implemented in a fully incremental fashion, with per-round runtime depending only on d , namely $\tilde{O}(d)$ for the first-order methods (EG, Soft-Bayes) and $O(d^2)$ for ONS.

The next improvement was achieved in [29]. In their algorithm called Ada-BARRONS, the authors managed to reach a polylogarithmic in T regret without a gradient boundedness assumption. This was done by combining ONS with a logarithmic barrier regularizer, a clever strategy of adaptive stepsize selection, and the use of self-concordance in the analysis. However, the regret remained suboptimal, scaling as $O(d^2 \log^4(T))$. In addition, the per-round runtime deteriorated to $\tilde{O}(d^{2.5} T)$, thus becoming T -dependent. Very recently, two competing works [49, 34] achieved a T -independent per-round runtime while preserving the $\tilde{O}(d^2)$ regret guarantee of Ada-BARRONS. In both cases, the crucial step was to combine Ada-BARRONS with an appropriate scheme of adaptive restarts.

Despite all these efforts and a plethora of computationally feasible algorithms resulting from them, Universal Portfolios has remained the only algorithm so far with an optimal $O(d \log(T))$

⁴Namely, finding $\operatorname{argmin}_{w \in \Delta_d} \|w - \hat{w}\|_{\mathbf{H}}^2$ for a given $\hat{w} \in \mathbb{R}^d$, where $\mathbf{H} \succ 0$. Choletsky’s decomposition translates this to a second-order cone program which can be solved in $\tilde{O}(d^3)$ via interior-point methods; see [37, Section 4.3.6].

regret. The challenge of providing a practical algorithm with a similar regret guarantee is well known. The interest to it was reignited in [46], where it was conjectured that Follow-The-Leader regularized with a logarithmic barrier (LB-FTRL)—an algorithm coinciding with VB-FTRL without a volumetric regularizer, i.e. with $\mu = 0$ —is regret-optimal. Since such an algorithm could be implemented with $\tilde{O}(d^2T)$ per-round runtime via Newton’s method, this would result in a regret-optimal and computationally feasible procedure. Yet, [49] recently disproved this conjecture: as it turns out, the regret of LB-FTRL can be of order $2^d \log(T) \log(\log(T))$ whenever $T > \text{poly}(2^d)$. In addition to the arguments we give in Section 2, this negative result demonstrates that adding a volumetric-barrier regularization term is crucial for obtaining a regret-optimal FTRL-type strategy.

Connections with cutting-plane and interior point methods. An interesting byproduct of our work is shedding new light on the *volumetric barrier of a polytope*, an object first studied by P. Vaidya in his seminal work [45] in the context of cutting-plane methods for convex optimization. Vaidya’s goal was to supercede the ellipsoid method [44, 24, 36] as the state-of-the-art cutting plane method, and he used the volumetric barrier as the key component of his algorithm. Once the general theory of interior-point methods (IPM) had been developed by Nesterov and Nemirovskii [38], the volumetric barrier was used by K. Anstreicher [3] as a self-concordant barrier for a polytope, replacing the logarithmic barrier in this role; this resulted in a faster-converging IPM for linear programming. Remarkably, in both these scenarios further improvements were obtained using the “hybrid” volumetric-logarithmic barrier (see [45, 3]), which mimics the potential used in VB-FTRL (see Section 3) except for a different choice of regularization parameter μ (namely, [3] prescribes $\mu = \frac{T+d-1}{d-1}$). To our best knowledge, the volumetric and hybrid barriers have never been used in online learning. Their emergence in this context—and in such a vital role—came as a suprise to us, and then challenged us to derive VB-FTRL as an approximation of Universal Portfolios (see Section 2.2).

2 Proposed algorithm and its connection with Universal Portfolios

In this section, we present VB-FTRL and derive it as an approximation of Universal Portfolios [12].

First of all, let us establish the notation and briefly recap the setup of online portfolio selection.

Notation. We let $\bar{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$. Assuming $d > 1$, we denote by \mathbb{R}_+^d the nonnegative orthant in \mathbb{R}^d , by Δ_d the standard simplex in \mathbb{R}^d , and by $\text{ri}(\Delta_d)$ the relative interior of Δ_d . We let e_i be the i^{th} canonical basis vector in \mathbb{R}^d , and $x[i]$ be the i^{th} entry of x . $\mathbf{1}_d$ be the vector of all ones in \mathbb{R}^d , and \mathbf{I}_d is the $d \times d$ identity matrix; we also write \mathbf{I} when the dimension is clear from the context. $[n]$ be the set of the first n positive integers; $[0]$ is the empty set. We use standard asymptotic notation: given two functions $f, g > 0$ of the same argument $a > 1$, notation $g = O(f)$ tells that $g(a) \leq cf(a)$ for some generic constant $c > 0$ and all a in the common domain of f and g ; $g = \tilde{O}(f)$ means that $g(a) \leq cf(a) \log^c(ea)$. We let $\|\cdot\|_p$, $p \geq 1$, be the ℓ_p -norm on \mathbb{R}^n . We let $\mathbb{S}^n, \mathbb{S}_+^n, \mathbb{S}_{++}^n$ be the sets of $n \times n$ symmetric, positive-semidefinite, and positive-definite matrices, and we use boldfaced capital letters for such matrices. For $\mathbf{M} \in \mathbb{S}_+^n$ and $u, v \in \mathbb{R}^n$, we let $\|u\|_{\mathbf{M}} := \sqrt{u^\top \mathbf{M} u}$ and $\langle u, v \rangle_{\mathbf{M}} := u^\top \mathbf{M} v$. We let $\|M\|$ the operator norm of a matrix M . We use the “Matlab” matrix notation: $[A_1, A_2]$ (resp. $[A_1; A_2]$) is the horizontal (resp. vertical) concatenation of two matrices A_1, A_2 with compatible dimensions.

Affine reparametrization. Assuming that $d > 1$,⁵ simplex $\Delta_d \subset \mathbb{R}^d$ has empty interior, which prevents us from using the machinery of self-concordant functions. To circumvent this issue, define

$$\mathbf{\Delta}_{d-1} := \{v \in \mathbb{R}_+^{d-1} : \mathbf{1}_{d-1}^\top v \leq 1\}, \quad (7)$$

a “solid” simplex in \mathbb{R}^{d-1} , and note that the affine mapping $v \mapsto Av + e_d$ with a $d \times (d-1)$ matrix

$$A := \begin{bmatrix} \mathbf{I}_{d-1} \\ -\mathbf{1}_{d-1}^\top \end{bmatrix} \quad (8)$$

bijectionally maps $\mathbf{\Delta}_{d-1}$ onto Δ_d , and \mathbb{R}^{d-1} onto the affine span of Δ_d , i.e. onto the hyperplane

$$\mathbb{A}_d := \{w \in \mathbb{R}^d : w^\top \mathbf{1}_d = 1\}. \quad (9)$$

As such, restricting a function $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ to \mathbb{A}_d and reparametrizing results in $f^\mathbf{1} : \mathbb{R}^{d-1} \rightarrow \overline{\mathbb{R}}$,

$$f^\mathbf{1}(v) := f(Av + e_d), \quad (10)$$

with the gradient and Hessian as follows:

$$\nabla f^\mathbf{1}(v) = A^\top \nabla f(Av + e_d) \quad \text{and} \quad \nabla^2 f^\mathbf{1}(v) = A^\top \nabla^2 f(Av + e_d) A. \quad (11)$$

Since A has full column rank, $f^\mathbf{1}$ is strictly convex (on \mathbb{R}^{d-1}) whenever f is strictly convex on \mathbb{A}_d . Finally, one can easily check that the mapping $v \mapsto Av + e_d$ of \mathbb{R}^{d-1} to \mathbb{A}_d has the inverse given by

$$w \mapsto A^+(w - e_d) \quad (12)$$

where $A^+ = (A^\top A)^{-1} A^\top$, $A^+ \in \mathbb{R}^{(d-1) \times d}$, is the left pseudoinverse of A , i.e. such that $A^+ A = \mathbf{I}_{d-1}$.

Summary of the setup. Online portfolio selection is a game between the *learner* (i.e. the trader) and her *adversary* (the market) played over $T \geq 1$ rounds according to the following protocol:

In each round $t \in [T]$:

1. The learner selects a new portfolio, i.e. a distribution $w_t \in \Delta_d$.
2. The adversary observes w_t and selects a new vector of asset returns $x_t \in \mathbb{R}_+^d$.
3. The learner observes x_t and suffers the loss $\ell_t(w_t)$ where $\ell_t(w) := -\log(x_t^\top w)$ for $w \in \Delta_d$.

Formally, an online portfolio selection *algorithm* or *strategy* is a sequence of mappings $(\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_{T-1})$ where \mathcal{S}_{t-1} maps the history $(w_\tau, x_\tau)_{\tau \in [t-1]}$ on Δ_d , that is $w_t = \mathcal{S}_{t-1}(w_1, x_1, \dots, w_{t-1}, x_{t-1})$ for $t > 1$ and $w_1 = \mathcal{S}_0(\emptyset)$. Pitting it against an adversary, as per the above protocol, results in two sequences

$$w_{1:T} := (w_1, \dots, w_T), \quad x_{1:T} := (x_1, \dots, x_T). \quad (13)$$

As previously discussed, we quantify the performance of an algorithm that produced a sequence of portfolios $w_{1:T}$ for the market realization $x_{1:T}$, in terms of its *regret* against the best CRP for $x_{1:T}$:

$$\mathcal{R}_T(w_{1:T} | x_{1:T}) := \sum_{t \in [T]} \ell_t(w_t) - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w).$$

⁵The case $d = 1$ is trivial: $w_t \equiv 1$ is a unique possible strategy. As such, from now on we tacitly assume that $d > 1$.

Following [12] and subsequent works, we aim at constructing an algorithm for which $\mathcal{R}_T(w_{1:T}|x_{1:T})$ admits an upper bound of the form (5) *uniformly* over all choices of $x_1, \dots, x_T \in \mathbb{R}_+^d$. Note that such a bound would imply that at any round $t \in [T]$ the adversary is allowed to know the learner’s *algorithm*—i.e. the whole sequence $(\mathcal{S}_0, \dots, \mathcal{S}_{T-1})$ rather than just the portfolios $w_{1:t}$ selected so far.

We are now ready to present our proposed algorithm and a regret guarantee that it admits.

2.1 Algorithm and main result

Our algorithm relies on some auxiliary functions. Let $\lambda, \mu > 0$ be two regularization parameters whose values will be specified a bit later. The *logarithmic barrier* $R : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ of \mathbb{R}_+^d is defined by

$$R(w) := \begin{cases} -\sum_{i \in [d]} \log(w[i]) & \text{if } w \in \mathbb{R}_{++}^d, \\ +\infty & \text{otherwise.} \end{cases} \quad (14)$$

Note that, $R(w)$ is proper, lower semicontinuous, and strictly convex. Moreover, $R^\mathbb{1}(v)$, cf. (7)–(10), is a barrier on Δ_{d-1} —in other words, $\text{dom}(R^\mathbb{1}) = \text{int}(\Delta_{d-1})$, and $R^\mathbb{1}(v)$ diverges as v approaches the boundary of Δ_{d-1} . Now, let $L_t(w)$ be the observed cumulative loss regularized by $R(w)$, namely

$$L_t(w) := \sum_{\tau \in [t]} \ell_\tau(w) + \lambda R(w) \quad (15)$$

for $\lambda > 0$ to be specified later; note that $L_0(w) = \lambda R(w)$. Define the *volumetric barrier* $V_t : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$:

$$V_t(w) := \begin{cases} \frac{1}{2} \log \det(A^\top \nabla^2 L_t(w) A) & \text{if } w \in \mathbb{R}_{++}^d, \\ +\infty & \text{otherwise.} \end{cases} \quad (16)$$

Finally, for $\mu > 0$ that shall also be specified later, define the potential function $P_t : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ by

$$P_t(w) := L_t(w) + \mu V_t(w). \quad (17)$$

Our proposed algorithm, called VB-FTRL, amounts to iteratively minimizing this potential, that is:

$$w_t = \underset{w \in \Delta_d}{\text{argmin}} P_{t-1}(w). \quad (\text{VB-FTRL})$$

Correctness of this update—i.e. that w_t defined above exists and is unique—follows from the strict convexity of V_t for all $t \in [T]$ and $\lambda, \mu > 0$. For $\lambda = \mu = 1$ this is well known [45, 3], and the general case is analogous (see Appendix C). In fact, direct differentiation allows to obtain explicit formulae for the gradient and Hessian of $P_t(w)$; following [3], one can then show that $P_t^\mathbb{1}(v)$ is a self-concordant function on Δ_{d-1} . (We provide these conceptually straightforward, but tedious calculations in Appendix C, and give some background on self-concordant functions in Appendix A.3.) Such functions are known as the “natural” class of objectives for Newton’s method, on which this method admits affine-invariant global convergence guarantees (see [37, Chapter 4]). This naturally leads to an algorithmic implementation of VB-FTRL via Newton’s method; we shall present this implementation and discuss its tractability in Section 4. Postponing further discussion of the computational matters to Section 4, let us now state and discuss a regret guarantee for VB-FTRL.

Theorem 2.1 (Regret of VB-FTRL). *For any $T \in \mathbb{N}$ and market realization $x_{1:T}$, one has that*

$$\mathcal{R}_T(w_{1:T}|x_{1:T}) \leq (\lambda + 2\mu)(d - 1) \log(T + \lambda d)$$

for the sequence $w_{1:T}$ produced by (VB-FTRL) with λ, μ satisfying $\frac{1}{\lambda}(1 + \frac{2\mu}{\lambda})^2 \leq \min\left\{\frac{1}{4}, \frac{5\mu}{8(1+\lambda)}\right\}$ and $\lambda \geq 2e$. In particular, with $\lambda = 16$ and $\mu = 7$ we guarantee that

$$\mathcal{R}_T(w_{1:T}|x_{1:T}) \leq 30(d - 1) \log(T + 16d).$$

On the one hand, the regret bound for VB-FTRL matches the bound for Universal Portfolios, albeit with a worse constant factor (cf. (5)), and also matches the worst-case lower bound from [11]. In other words, the regret optimality of Universal Portfolios is preserved in the proposed algorithm.

On the other hand, the update in VB-FTRL reduces to solving a *convex optimization* problem, and thus can be efficiently implemented; as such, we address the challenge put forward in Section 1. This is in stark contrast with Universal Portfolios where one has to compute a multivariate integral in each round. In fact, the update in Universal Portfolios can be seen as minimizing a certain convex functional on the space of *probability measures supported on Δ_d* , which is way more challenging, from an algorithmic standpoint, than convex optimization on Δ_d . In Section 2.2 we discuss this point in more detail, and shed light on the connection between VB-FTRL and Universal Portfolios.

Finally, putting $\mu = 0$ in (VB-FTRL) we recover Follow-The-Leader regularized with a logarithmic barrier (LB-FTRL), the algorithm conjectured to have an $O(d \log(T))$ regret in [46]. As we have already observed in Section 1.1, the recent refutation of this conjecture in [49] indicates that time-dependent regularization is *instrumental* in achieving regret of the right order with an FTRL-type strategy. The regret analysis of VB-FTRL in Section 3 demonstrates this on a technical level.

2.2 VB-FTRL as an approximation of Universal Portfolios

In this section, our goal is to show that (VB-FTRL) naturally arises as an approximation of the Universal Portfolios update (3); moreover: informally speaking, the accuracy of this approximation can be controlled in natural terms. Although arguably quite natural, the argument about to be presented is a rather delicate one, and the corresponding accuracy bound depends on some auxiliary objects that have to be introduced first. For the sake of exposition, we shall split this argument into three logical steps that are briefly outlined below, and defer some technical details to Appendix B.

- 1^o. We first incorporate two parameters (λ, μ) into (3), so that “vanilla” Cover’s update corresponds to $\lambda = 1, \mu = 0$. We then establish Proposition 2.1 stating that the generalized strategy is optimal for any (constant) λ and μ . Besides, we observe that the corresponding distribution ϕ_t admits a variational formulation as a minimizer of the Gibbs’ potential $F_{t-1}[\cdot]$,

$$\phi_t \in \underset{\phi \in \text{Supp}(\Delta_d)}{\text{argmin}} F_{t-1}[\phi]; \tag{18}$$

here $\text{Supp}(\Delta_d)$ is the set of distributions supported on Δ_d , and $F_{t-1}[\cdot]$ shall be defined in (23).

2^o. Next we focus on minimization problem (18). Using self-concordance [37], we show that

$$c \operatorname{tr}(\nabla^2 L_{t-1}(\hat{w}) \operatorname{Cov}[\phi]) - \mu \operatorname{Ent}[\phi] \leq F_{t-1}[\phi] - L_{t-1}(\hat{w}) \leq C \operatorname{tr}(\nabla^2 L_{t-1}(\hat{w}) \operatorname{Cov}[\phi]) - \mu \operatorname{Ent}[\phi] \quad (19)$$

for any distribution ϕ with expectation $\hat{w} \in \Delta_d$, covariance $\operatorname{Cov}[\phi]$, differential entropy $\operatorname{Ent}[\phi]$, and supported on rescaled Dikin ellipsoid $\mathcal{E}_{t-1,1/2}(\hat{w}) := \{w \in \mathbb{A}_d : \|w - \hat{w}\|_{\nabla^2 L_{t-1}(\hat{w})} \leq 1/2\}$. It is well known that $\mathcal{E}_{t-1,1}(w)$ is contained in Δ_d for any $w \in \Delta_d$ [37]; as such, a natural idea is to approximate ϕ_t by minimizing the upper bound for $F_{t-1}[\phi]$, as per (19), over such distributions. This relaxation is indeed reasonable: the corresponding minimizer $\bar{\phi}_t$ satisfies

$$F_{t-1}[\bar{\phi}_t] \leq \min_{\phi \in \operatorname{Supp}(\Delta_d)} F_{t-1}[\phi] + O(\mu d \log(T + \lambda d)). \quad (20)$$

That is, $F_{t-1}[\cdot]$ is minimized up to an error of the order of the desired regret. This result, rigorously formulated as Proposition 2.2, is proved via self-concordance machinery and some volume estimates, the proof being rather technical; we believe it to be of independent interest.

3^o. We then focus on the relaxation of (18) described in step 2^o, i.e. replacing $F_{t-1}[\phi]$ with its upper bound from (19), and imposing the constraint $\phi \in \operatorname{Supp}(\mathcal{E}_{t-1,c}(\hat{w}))$ for some $c \leq 1/2$. We relax this “hard” support constraint into a “soft” covariance constraint, namely

$$\phi \in \operatorname{Supp}(\mathbb{A}_d), \quad \operatorname{Cov}[\phi] \preceq c^2 L_{t-1}(\hat{w})^{-1}.$$

As it happens, the minimum is attained on the \mathbb{A}_d -marginal of the Gaussian distribution with covariance $\nabla^2 L_{t-1}(\hat{w})^{-1}$ and expectation w_t —precisely as in (VB-FTRL). The correspondence between (VB-FTRL) and (generalized) Cover’s update is thus specified. Moreover, we show (see Proposition 2.3) that the truncation of this Gaussian distribution to the appropriate Dikin ellipsoid around w_t is an approximate minimizer of $F_{t-1}[\cdot]$ with accuracy $O(\mu d \log(ed))$ —in other words, the Gaussian relaxation does not lead to a further loss of accuracy beyond (20). In our view, this result—just as (20)—may find uses beyond the context of portfolio selection.

We are now about to implement these steps. Before doing so, let us make one remark. One might expect that the above approximation result—the one representing (VB-FTRL) update as an approximation of generalized Universal Portfolios with the “right” accuracy of order $O(d \log(T))$ —would directly lead to a shorter and/or simpler proof of Theorem 2.1 than the one to be presented in Section 3, via some formal reduction to the regret bound for (λ, μ) -Universal Portfolios in Proposition 2.1. However, we were unable to find such an alternative proof. In fact, the actual proof of Theorem 2.1 in Section 3 has very little to do—if anything at all—with the proofs of Propositions 2.2–2.3—in particular, it is non-probabilistic and does not use the variational formulation (18).

Step 1^o : (λ, μ) -generalized Universal Portfolios update and its variational formulation.

Recall that Cover’s original method (3)–(4) amounts to playing the expectation over the density proportional to $\exp(-\sum_{\tau \in [t-1]} \ell_\tau(w))$, $w \in \Delta_d$, at round t . Now, consider the generalized strategy

$$w_t = \mathbb{E}_{w \sim \phi_t}[w] \quad \text{where} \quad \phi_t(w) \propto \exp\left(-\frac{1}{\mu} L_{t-1}(w)\right), \quad w \in \Delta_d, \quad (21)$$

where we introduced two parameters $\lambda \geq 0$ and $\mu \geq 1$, so that (4) corresponds to $\lambda = 0$, $\mu = 1$. Here, parameter μ plays the role of temperature: the larger it is, the closer ϕ_t is to the *uniform*

distribution on Δ_d . Meanwhile, λ biases ϕ_t towards the Dirac measure on $\frac{1}{d}\mathbb{1}_d$, the uniform portfolio. We first observe that the regret optimality of Cover's algorithm extends to the whole range of λ, μ .

Proposition 2.1. *For any $T \in \mathbb{N}$ and any market realization $x_{1:T}$, the sequence of portfolios $w_{1:T}$ constructed by (λ, μ) -generalized Universal Portfolios (21) with parameters $\lambda \geq 0, \mu \geq 1$ satisfies:*

$$\mathcal{R}_T(w_{1:T}|x_{1:T}) \leq \mu(d-1) \log(T+1) + \lambda(d-1) \log\left(4e \frac{T+\lambda d}{\lambda d}\right) + \mu - \lambda \log(d).$$

The argument, given in Appendix B, extends the proof of [20, Theorem 7] to allow for $\lambda > 0$. The key observation is that $\exp(-\frac{1}{\mu}\ell_t(w)) = (x_t^\top w)^{1/\mu}$ is a concave function whenever $\mu \geq 1$. By Jensen's inequality, this leads, through a telescoping argument, to the bound as follows:

$$\sum_{t \in [T]} \ell_t(w) \leq -\mu \log \left[\int_{\Delta_d} \exp\left(-\frac{1}{\mu}L_T(w)\right) dw \right] + \mu \log \left[\int_{\Delta_d} \exp\left(-\frac{\lambda}{\mu}R(w)\right) dw \right]. \quad (22)$$

The right-hand side is then related to $\min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w)$ through the rescaling technique from [7].

For what is to follow, it is crucial to make a couple of observations. First, ϕ_t , as defined in (21), is the unique optimal solution to the following optimization problem (see e.g. [47, Lemma 4.10]):

$$\min_{\phi \in \text{Supp}(\Delta_d)} F_{t-1}[\phi] := \mathbb{E}_{w \sim \phi}[L_{t-1}(w)] - \mu \text{Ent}[\phi] \quad (23)$$

where $\text{Supp}(\Delta_d)$ is the set of probability densities supported on Δ_d , and $\text{Ent}[\phi] := \mathbb{E}_{w \sim \phi}[-\log \phi(w)]$ is the differential entropy of ϕ . Moreover, $F_{t-1}[\phi_t]$ is the (rescaled) negative log-partition function:

$$F_{t-1}[\phi_t] = -\mu \log \int_{\Delta_d} \exp\left(-\frac{1}{\mu}L_{t-1}(w)\right) dw.$$

In particular, we recognize the first term in the right-hand side of (22) as nothing else but $F_T[\phi_{T+1}]$.

Step 2^o : restriction to the Dikin ellipsoid. Fix $\hat{w} \in \text{ri}(\Delta_d)$, and for $r \geq 0$ and $t \in [T]$ define

$$\mathcal{E}_{t-1,r}(\hat{w}) := \{w \in \mathbb{A}_d : \|w - \hat{w}\|_{\nabla^2 L_{t-1}(\hat{w})} < r\}, \quad (24)$$

the r -Dikin ellipsoid [37] of $L_{t-1}(\cdot)$ at \hat{w} restricted to hyperplane \mathbb{A}_d , cf. (9). It can be routinely checked (see Appendix A.3) that when $\lambda \geq 1$, the function $L_{t-1}^{\lambda}(v) = L_{t-1}(Av + e_d)$ on \mathbb{R}^{d-1} , cf. (10), is self-concordant with domain $\text{int}(\Delta_{d-1})$. By [37, Theorems 4.1.5], the ‘‘unit’’ ellipse $\mathcal{E}_{t-1,1}(\hat{w})$ is contained in $\text{ri}(\Delta_d)$. Moreover, on smaller ellipses one can control the remainder of the quadratic expansion of L_t : from [37, Theorems 4.1.6–4.1.8] it follows (see Lemma A.8) that

$$\frac{1}{5} \|w - \hat{w}\|_{\nabla^2 L_{t-1}(\hat{w})}^2 \leq L_{t-1}(w) - L_{t-1}(\hat{w}) - \nabla L_{t-1}(\hat{w})^\top (w - \hat{w}) \leq \frac{4}{5} \|w - \hat{w}\|_{\nabla^2 L_{t-1}(\hat{w})}^2 \quad (25)$$

for all $w \in \mathcal{E}_{t-1,1/2}(\hat{w})$.

Hence, any distribution ϕ with mean $\mathbb{E}_\phi[w] = \hat{w} \in \text{ri}(\Delta_d)$ and supported in $\mathcal{E}_{t-1,1/2}(\hat{w})$, satisfies

$$\frac{1}{5} \text{tr}(\nabla^2 L_{t-1}(\hat{w}) \text{Cov}[\phi]) - \mu \text{Ent}[\phi] \leq F_{t-1}[\phi] - L_{t-1}(\hat{w}) \leq \frac{4}{5} \text{tr}(\nabla^2 L_{t-1}(\hat{w}) \text{Cov}[\phi]) - \mu \text{Ent}[\phi] \quad (26)$$

where $\text{Cov}[\phi]$ is the covariance matrix of ϕ . The upper bound suggests us to approximate (23) with

$$\min_{\substack{\hat{w} \in \Delta_d, \mathbb{E}_\phi[w] = \hat{w}, \\ \phi \in \text{Supp}(\mathcal{E}_{t-1, 1/4}(\hat{w}))}} \underbrace{L_{t-1}(\hat{w}) + \frac{4}{5} \text{tr}(\nabla^2 L_{t-1}(\hat{w}) \text{Cov}[\phi]) - \mu \text{Ent}[\phi]}_{:= \bar{F}_{t-1}[\phi]} \quad (27)$$

—in other words, to focus on distributions supported in a small Dikin ellipsoid around the expectation. Now: since $\mathcal{E}_{t-1, 1/4}(\hat{w}) \subset \Delta_d$, any optimal solution $\bar{\phi}_t$ to (27) is feasible in (23); hence (27) overestimates (23). Moreover, the following result gives the accuracy of this upper approximation:

Proposition 2.2. *When $\lambda \geq 1$, any optimal solution $\bar{\phi}_t$ to (27) satisfies the following inequalities:*

$$F_{t-1}[\bar{\phi}_t] \leq \bar{F}_{t-1}[\bar{\phi}_t] \leq \min_{\phi \in \text{Supp}(\Delta_d)} F_{t-1}[\phi] + 1.5\mu(d-1) \log(T + \lambda d) + 3.9\mu(d+1) + 0.1.$$

The proof of this result is relegated to Appendix B and proceeds in two steps. First, we show that

$$F_{t-1}[\bar{\phi}_t] \leq \bar{F}_{t-1}[\bar{\phi}_t] \leq \min_{\substack{\hat{w} \in \Delta_d, \mathbb{E}_\phi[w] = \hat{w}, \\ \phi \in \text{Supp}(\mathcal{E}_{t-1, 1/2}(\hat{w}))}} F_{t-1}[\phi] + \mu(d-1) \log(2).$$

Here the first inequality is trivial, and the second one follows by comparing the covariance and entropy for the random variables $w \sim \phi$ and $\hat{w} + \frac{1}{2}(w - \hat{w})$. The second step consists in showing that (λ, μ) -generalized Universal Portfolios distribution ϕ_t , as defined in (21), can be truncated to the Dikin ellipsoid around its mode at the expense of an $O(\mu(d-1) \log(T + \lambda d))$ increase in F_{t-1} . This is done by expressing $F_{t-1}[\phi_t]$ as the log-partition function, controlling the integrand via self-concordance, and lower-bounding the volume of the Dikin ellipsoid by arguing that the mode of ϕ_t cannot be too close to the relative boundary of Δ_d thanks to the logarithmic-barrier regularizer.

Step 3^o : reduction to the Gaussian distribution. Departing from (27), we note that its objective $\bar{F}_{t-1}[\phi]$ depends on ϕ only through the mean, covariance, and differential entropy of ϕ . However, the support constraint, namely $\phi \in \text{Supp}(\mathcal{E}_{t-1, 1/4}(\hat{w}))$, interferes with higher moments. A natural idea is then to replace the “hard” support constraint with a “soft” covariance one, namely

$$\phi \in \text{Supp}(\mathbb{A}_d), \quad \text{Cov}[\phi] \preceq \frac{1}{16} \nabla^2 L_{t-1}(\hat{w})^{-1}. \quad (28)$$

This results in a relaxation of (27). Indeed, $\mathbb{E}_{w \sim \phi}[w] = \hat{w} \in \Delta_d$ and $\phi \in \text{Supp}(\mathcal{E}_{t-1, 1/4}(\hat{w}))$ together imply that $\mathbb{1}_d^\top (w - \hat{w}) = 0$ (ϕ -a.s.), that is $\phi \in \text{Supp}(\mathbb{A}_d)$; on the other hand, by Jensen’s inequality

$$\left\| \nabla^2 L_{t-1}(\hat{w})^{1/2} \text{Cov}[\phi] \nabla^2 L_{t-1}(\hat{w})^{1/2} \right\| \leq \mathbb{E}_{w \sim \phi} \left[\|w - \hat{w}\|_{\nabla^2 L_{t-1}(\hat{w})}^2 \right] \leq \frac{1}{16}. \quad (29)$$

We can now satisfy the support constraint in (28) “automatically” by reformulating it as follows:

$$w \sim \phi \text{ is such that } w = Av + e_d \text{ for } v \sim \phi^\mathbb{1} \in \text{Supp}(\mathbb{R}^{d-1}) \quad (30)$$

where the mean $\hat{v} \in \mathbf{\Delta}_{d-1}$, covariance, and entropy of $\phi^\mathbb{1}$ are related to those of ϕ according to

$$\hat{w} = A\hat{v} + e_d, \quad \text{Cov}[\phi] = A \text{Cov}[\phi^\mathbb{1}] A^\top, \quad \text{Ent}[\phi] = \text{Ent}[\phi^\mathbb{1}] - \frac{1}{2} \log \det(A^\top A)$$

where $\det(A^\top A) = d$. The matrix constraint in (28) translates to $\text{Cov}[\phi^1] \preceq \frac{1}{16}(A^\top \nabla^2 L_{t-1}(\hat{w})A)^{-1}$ by simple linear algebra. As such, in terms of \hat{w} , $\text{Cov}[\phi^1]$, and $\text{Ent}[\phi^1]$, the objective in (27) becomes

$$\bar{F}_{t-1}[\phi] = L_{t-1}(\hat{w}) + \frac{4}{5} \text{tr} \left[\text{Cov}[\phi^1] A^\top \nabla^2 L_{t-1}(\hat{w})A \right] - \mu \text{Ent}[\phi^1]. \quad (31)$$

Clearly, $\mathcal{N}(\hat{v}, \Sigma)$ maximizes $\text{Ent}[\phi^1]$ given the constraints $\mathbb{E}_{\phi^1}[v] = \hat{v}$ and $\text{Cov}[\phi^1] = \Sigma$, with the value $\text{Ent}[\mathcal{N}(\hat{v}, \Sigma)] = \frac{1}{2} \log \det(\Sigma) + c_d$ where c_d depends only on d . This gives a relaxation of (27),

$$\min_{\hat{w} \in \Delta_d, \quad 0 \preceq \Sigma \preceq \frac{1}{16}(A^\top \nabla^2 L_{t-1}(\hat{w})A)^{-1}} L_{t-1}(\hat{w}) + \frac{4}{5} \text{tr} \left[\Sigma A^\top \nabla^2 L_{t-1}(\hat{w})A \right] - \frac{\mu}{2} \log \det(\Sigma), \quad (32)$$

that depends on ϕ only through the mean and covariance. Finally, in Appendix B.4 we show that, when $\mu > 0.1$, the minimum over Σ is attained at $\hat{\Sigma} = \frac{1}{16}(A^\top \nabla^2 L_{t-1}(\hat{w})A)^{-1}$, and (32) reduces to

$$\min_{\hat{w} \in \Delta_d} L_{t-1}(\hat{w}) + \frac{\mu}{2} \log \det(A^\top \nabla^2 L_{t-1}(\hat{w})A). \quad (33)$$

This is nothing else but (VB-FTRL). To summarize, (VB-FTRL) appears after relaxing the ‘‘hard’’ support constraint in (27), so that the optimum is attained on a Gaussian distribution. Moreover: the result below shows that, essentially, this relaxation comes at no extra cost (cf. Proposition 2.2).

Proposition 2.3. *Assume that $\mu > 0.1$. Let $\bar{\phi}_t$ be optimal in (27), w_t be as in (VB-FTRL), $\Sigma_t = \frac{1}{16}(A^\top \nabla^2 L_{t-1}(w_t)A)^{-1}$, and let g_t^{trc} be the truncation of $\mathcal{N}(w_t, A\Sigma_t A^\top)$ to $\mathcal{E}_{t-1, 1/2}(w_t)$. Then*

$$F_{t-1}[g_t^{\text{trc}}] \leq \bar{F}_{t-1}[\bar{\phi}_t] + 0.5\mu(d-1) \log(350d) + 0.01\mu.$$

The proof of this result relies on the feasibility of $\bar{\phi}_t$ in (32) and an estimate of the decrease of entropy when truncating a Gaussian distribution to its covariance ellipsoid (see Lemma B.2).

Finally, by combining Propositions 2.2 and 2.3 we establish a quantitative correspondence between (VB-FTRL) and generalized Cover’s update in its variational form (18). Namely, portfolio w_t computed in (VB-FTRL) generates a distribution—precisely, $\mathcal{N}(w_t, \frac{1}{16}A(A^\top \nabla^2 L_{t-1}(w_t)^{-1}A)^{-1}A^\top)$ truncated to the Dikin ellipsoid $\mathcal{E}_{t-1, 1/2}(w_t)$ —that solves (18) up to the accuracy of $O(\mu d \log(T))$.

3 Regret analysis

In this section, we start by giving an outline of the proof of Theorem 2.1 and discussing the role of volumetric regularization. After that, we are going to present the full proof of the theorem.

3.1 Overview of the proof and the role of volumetric regularization

Key decomposition. We shall specify the values of λ, μ in the later stages of the proof; for now we consider general $\lambda, \mu > 0$. Let us decompose the regret of the sequence generated in (VB-FTRL) through telescoping:

$$\begin{aligned} \mathcal{R}_T(w_{1:T}|x_{1:T}) &\leq \sum_{t \in [T]} \left[\ell_t(w_t) + \min_{w \in \Delta_d} P_{t-1}(w) - \min_{w \in \Delta_d} P_t(w) \right] - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w) + \min_{w \in \Delta_d} P_T(w) \\ &= \sum_{t \in [T]} \left[P_t(w_t) + \mu V_{t-1}(w_t) - \mu V_t(w_t) - P_t(w_{t+1}) \right] - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w) + \min_{w \in \Delta_d} P_T(w). \end{aligned} \quad (34)$$

Here the inequality holds since $P_0(w) = \lambda R(w) \geq 0$ on Δ_d , cf. (14). In other words, we have that

$$\mathcal{R}_T(w_{1:T}|x_{1:T}) \leq \text{Bias}_T + \sum_{t \in [T]} [\text{Miss}_t + \text{Gain}_t] \quad (35)$$

where the three kinds of differences $\text{Bias}_T, \text{Miss}_t, \text{Gain}_t$ are defined as follows:

$$\text{Miss}_t := P_t(w_t) - P_t(w_{t+1}), \quad \text{Gain}_t := \mu[V_{t-1}(w_t) - V_t(w_t)], \quad \text{Bias}_T := \min_{w \in \Delta_d} P_T(w) - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w).$$

Here, Bias_T is the “bias” term arising since (VB-FTRL) minimizes $P_{t-1}(w)$ instead of $\sum_{\tau \in [t-1]} \ell_\tau(w)$, the observed sum of losses—in other words, due to our use of regularization “per se.” Miss_t is the “lagged” term that arises since we only “follow” the (regularized) leader, i.e. minimize P_{t-1} instead of P_t . Finally, Gain_t appears due to our use of a *time-varying* volumetric regularizer. Since P_{t-1} in (VB-FTRL) includes V_{t-1} rather than V_t , we expect Gain_t to be *negative* and *decrease* the regret.

Big picture. Our proof proceed as follows. On the one hand, we shall bound the “bias” term as

$$\text{Bias}_T \leq (\lambda + 2\mu)(d - 1) \log(T + \lambda d) \quad (36)$$

under the assumption that $\lambda \geq 2e$. This is not too involved, and can be done via a combination of standard techniques. Way more interesting is another step, in which we show that the inequalities

$$\text{Miss}_t + \text{Gain}_t \leq 0, \quad \forall t \in [T], \quad (37)$$

hold when λ, μ are as in the premise of the theorem. The result then follows by combining (35)–(37).

Volumetric regularization mechanism. First, convexity of $\log \det(\cdot)$ on \mathbb{S}_+^n implies that

$$\text{Gain}_t \leq -\frac{\mu}{2} \pi_t(w_t) \quad \text{with} \quad \pi_t(w) := \left\| A^\top \nabla \ell_\tau(w) \right\|_{(A^\top \nabla^2 L_t(w) A)^{-1}}^2, \quad (38)$$

see Lemma A.4. Under the affine reparametrization in (10), $\pi_t(w_t)$ becomes the leverage score of x_t at w_t , cf. (11). On the other hand, Miss_t is the suboptimality gap of w_t in terms of P_t or $P_t^\mathbb{1}$:

$$\text{Miss}_t = P_t(w_t) - \min_{w \in \Delta_d} P_t(w) = P_t^\mathbb{1}(v_t) - \min_{w \in \Delta_d} P_t^\mathbb{1}(v).$$

Now, self-concordance of $P_t^\mathbb{1}$ allows to bound Miss_t in terms of the squared Newton decrement:

$$\text{Miss}_t \leq C \text{Decr}_t^2 \quad \text{with} \quad \text{Decr}_t = \left\| A^\top \nabla P_t(w_t) \right\|_{(A^\top \nabla^2 P_t(w_t) A)^{-1}}, \quad (39)$$

for some $C > 0$, once we have shown that $\text{Decr}_t^2 \leq c < 1$ (see Lemma A.7 or [37, Theorem 4.1.13]). Finally, the optimality condition for w_t implies that

$$A^\top \nabla P_t(w_t) = A^\top \nabla \ell_t(w_t) + \mu A^\top [\nabla V_t(w_t) - \nabla V_{t-1}(w_t)].$$

Departing from this identity and exploiting the algebraic structure of leverage scores, we manage to control the contribution of the term $\mu A^\top [\nabla V_t(w_t) - \nabla V_{t-1}(w_t)]$ to Decr_t^2 , as μ increases, by simultaneously increasing λ . More precisely, we show that $\text{Decr}_t^2 \leq \frac{1+\lambda}{\lambda} (1 + \frac{2\mu}{\lambda})^2 \pi_t(w_t)$, while $\pi_t(w_t) \leq \frac{1}{1+\lambda}$ by Lemma C.1. Thus, selecting λ appropriately we can enforce $\text{Decr}_t^2 \leq c < 1$ for *any* $\mu > 0$. But then, combining the bound on Decr_t^2 with (38)–(39) and selecting μ appropriately, we enforce (37).

Remark 3.1. In the actual proof, instead of Decr_t we focus on the Newton decrement $\underline{\text{Decr}}_t$ of

$$P_t^{\mathbb{1}}(v) = L_t^{\mathbb{1}}(v) + \mu[V_t^{\mathbb{1}}(v_t) + \nabla V_t^{\mathbb{1}}(v_t)^\top(v - v_t)],$$

a lower bound for $P_t^{\mathbb{1}}(v)$ tight at $v = v_t$. As it turns out, while $\underline{\text{Decr}}_t$ is somewhat larger than Decr_t , focusing on $\underline{\text{Decr}}_t$ allows to avoid explicit use of the self-concordance of $P_t^{\mathbb{1}}$, using self-concordance of $L_t^{\mathbb{1}}$ instead. (Of course, we use the differential properties of V_t , in particular the algebraic structure of its derivatives, when bounding $\underline{\text{Decr}}_t$.) The point is that $P_t^{\mathbb{1}}$ is only 21-self-concordant (see Lemma C.1), rather than 1-self-concordant as $L_t^{\mathbb{1}}$, so working with $\underline{\text{Decr}}_t, P_t^{\mathbb{1}}$ instead of $\text{Decr}_t, P_t^{\mathbb{1}}$ allows to reduce the constants. Meanwhile, in the proof of Theorem 4.1 on the regret of (VB-FTRL-qN), a procedure implementing (VB-FTRL), we have to use self-concordance of $P_t^{\mathbb{1}}$ directly because of a quasi-Newton method applied to $P_t^{\mathbb{1}}$; the resulting constant factor is an order of magnitude larger.

3.2 Proof of Theorem 2.1

1°. Our goal in this step is to establish (36). Let us further decompose the “bias” term as

$$\text{Bias}_T = \text{VolBias}_T + \text{LogBias}_T \quad (40)$$

where the two terms correspond to the two regularizers $V_T(w)$ and $R(w)$ respectively:

$$\text{VolBias}_T := \min_{w \in \Delta_d} P_T(w) - \min_{w \in \Delta_d} L_T(w), \quad \text{LogBias}_T := \min_{w \in \Delta_d} L_T(w) - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w).$$

Defining $w_T^* := \text{argmin}_{w \in \Delta_d} L_T(w)$ we then observe that

$$\begin{aligned} \text{VolBias}_T &\leq P_T(w_T^*) - L_T(w_T^*) = \mu V_T(w_T^*) \leq \frac{\mu}{2}(d-1) \log(\|A^\top \nabla^2 L_T(w_T^*) A\|) \\ &\leq \frac{\mu}{2}(d-1) \log(d \|\nabla^2 L_T(w_T^*)\|) \end{aligned} \quad (41)$$

where in the last step we used that $\|A\| = \sqrt{d}$. On the other hand, we can estimate $\|\nabla^2 L_T(w_T^*)\|$ using that the entries of w_T^* cannot be too small. Indeed, according to Lemma A.1 we have that

$$w_T^*[i] \geq \frac{\lambda}{T + \lambda d}, \quad \forall i \in [d] \quad (42)$$

—a variation of the well-known result for the logarithmic-barrier regularizer (see e.g. [29]). But then

$$\nabla^2 L_T(w_T^*) = \sum_{t \in [T]} \frac{x_t x_t^\top}{(x_t^\top w_T^*)^2} + \sum_{i \in [d]} \frac{\lambda e_i e_i^\top}{(w_T^*[i])^2} \preceq \frac{(T + \lambda d)^2}{\lambda^2} \left(\sum_{t \in [T]} \frac{x_t x_t^\top}{\|x_t\|_1^2} + \lambda \sum_{i \in [d]} e_i e_i^\top \right) \preceq \frac{(T + \lambda d)^3}{\lambda^2} I_d,$$

that is

$$\|\nabla^2 L_T(w_T^*)\| \leq \frac{(T + \lambda d)^3}{\lambda^2}. \quad (43)$$

Returning to (41) and using that $\lambda \geq 1$ to simplify the argument of the logarithm, we find that

$$\text{VolBias}_T \leq 2\mu(d-1) \log(T + \lambda d). \quad (44)$$

Now, in order to bound LogBias take any $w_T^o \in \text{Argmin}_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w)$ and define $w_{T,\alpha}^o \in \Delta_d$ by

$$w_{T,\alpha}^o = (1 - \alpha)w_T^o + \frac{\alpha}{d}\mathbf{1}_d \quad \text{with} \quad \alpha = \frac{\lambda(d-1)}{T + \lambda(d-1)}.$$

On the one hand, we have that

$$R(w_{T,\alpha}^o) \leq \max_{w \in \Delta_d} R\left((1 - \alpha)w + \frac{\alpha}{d}\mathbf{1}_d\right) = (d-1) \log\left(\frac{d}{\alpha}\right) \leq (d-1) \log\left(2\frac{T + \lambda d}{\lambda}\right)$$

where we used that a convex function with a compact domain is maximized at an extremal points, and also that $\frac{d}{d-1} \leq 2$. On the other hand, by Lemma A.2 proved in appendix, for any $w \in \Delta_d$

$$\sum_{t \in [T]} \ell_t\left((1 - \alpha)w + \frac{\alpha}{d}\mathbf{1}_d\right) - \ell_t(w) \leq \frac{\alpha}{1 - \alpha}T = \lambda(d-1),$$

whence $\sum_{t \in [T]} \ell_t(w_{T,\alpha}^o) - \ell_t(w_T^o) \leq \lambda(d-1)$. Combining these two results, we conclude that

$$\text{LogBias}_T \leq L_T(w_{T,\alpha}^o) - \sum_{t \in [T]} \ell_t(w_T^o) \leq \lambda R(w_{T,\alpha}^o) + \sum_{t \in [T]} \ell_t(w_{T,\alpha}^o) - \ell_t(w_T^o) \leq \lambda(d-1) \log\left(2e\frac{T + \lambda d}{\lambda}\right). \quad (45)$$

We arrive at the desired estimate, cf. (36), by plugging (44) and (45) into (40) and using that $\lambda \geq 2e$.

2°. Next we turn to proving (37). To streamline the exposition, we define $x_{-i} = e_i$ for $i \in [d]$, let τ vary over the extended index set $[t]^+ := [t] \cup \{-1, \dots, -d\}$, and define the abridged notation:

$$\begin{aligned} \mathbf{H}_t(w) &:= A^\top \nabla^2 L_t(w) A, & \nabla_\tau(w) &:= A^\top \nabla \ell_\tau(w), \\ \hat{\mathbf{H}}_t &:= \mathbf{H}_t(w_t), & \hat{\nabla}_\tau &:= \nabla_\tau(w_t). \end{aligned} \quad (46)$$

Moreover, with a slight abuse of notation we let

$$\lambda_\tau := \lambda^{\mathbf{1}\{\tau < 0\}}, \quad \tau \in [t]^+, \quad (47)$$

so that $L_t(w) = \sum_{\tau \in [t]^+} \lambda_\tau \ell_\tau(w)$ and, as a consequence,

$$\mathbf{H}_t(w) = \sum_{\tau \in [t]^+} \lambda_\tau \nabla_\tau(w) \nabla_\tau^\top(w), \quad \hat{\mathbf{H}}_t = \sum_{\tau \in [t]^+} \lambda_\tau \hat{\nabla}_\tau \hat{\nabla}_\tau^\top. \quad (48)$$

Finally, we define the Gram matrix $\mathbf{\Pi}(w) \in \mathbb{S}_+^{t+d}$ of the extended dataset, with entries given by

$$\pi_{\tau,\nu}(w) := \sqrt{\lambda_\tau \lambda_\nu} \langle \nabla_\tau(w), \nabla_\nu(w) \rangle_{\mathbf{H}_t(w)^{-1}}, \quad \forall \tau, \nu \in [t]^+, \quad (49)$$

rows/columns indexed on $[t]^+$, and dependence on t omitted for brevity. Also, let $\hat{\mathbf{\Pi}} := \mathbf{\Pi}(w_t)$, i.e.

$$\hat{\pi}_{\tau,\nu} := \sqrt{\lambda_\tau \lambda_\nu} \langle \hat{\nabla}_\tau, \hat{\nabla}_\nu \rangle_{\hat{\mathbf{H}}_t^{-1}}, \quad \forall \tau, \nu \in [t]^+. \quad (50)$$

For brevity, we refer to the diagonal entries with a single index: $\pi_\tau(w) := \pi_{\tau,\tau}(w)$ and $\hat{\pi}_\tau := \hat{\pi}_{\tau,\tau}$. When $w \in \Delta_d$, $\pi_\tau(w)$ is the squared leverage score for the reparametrized and $\sqrt{\lambda_\tau}$ -rescaled losses:

$$\pi_\tau(Av + e_d) = \lambda_\tau \|\nabla \ell_\tau^{\mathbf{1}}(v)\|_{\nabla^2 L_t^{\mathbf{1}}(v)^{-1}}^2.$$

In fact, $\mathbf{\Pi}(w)$ is a projection matrix, i.e. $\mathbf{\Pi}(w) = \mathbf{\Pi}(w)^2$; see [45]. We will use a weaker property:

$$\pi_\tau(w) = \sum_{\nu \in [t]^+} \pi_{\tau,\nu}(w)^2, \quad \forall \tau \in [t]^+. \quad (51)$$

For the paper to be self-contained, this and other properties of $\mathbf{\Pi}(w)$ are proved in Appendix C.

3^o. Now, in terms of the notation we have just introduced, for each term Gain_t we have:

$$\text{Gain}_t = \frac{\mu}{2} \log \left(\frac{\det \left(\hat{\mathbf{H}}_t - \hat{\nabla}_t \hat{\nabla}_t^\top \right)}{\det \left(\hat{\mathbf{H}}_t \right)} \right) \leq -\frac{\mu}{2} \text{tr} \left[\hat{\nabla}_t \hat{\nabla}_t^\top \hat{\mathbf{H}}_t^{-1} \right] = -\frac{\mu}{2} \hat{\pi}_t. \quad (52)$$

Here the inequality relies on the fact that $\log \det(\cdot)$ is a convex function on \mathbb{S}_+^n (see Lemma A.4). As such, to verify (37) it suffices to show that, for the values of λ, μ as in the premise of the theorem,

$$\text{Miss}_t \leq \frac{\mu}{2} \hat{\pi}_t. \quad (53)$$

To that end, let us define $\underline{P}_t : \mathbb{R}_{++}^d \rightarrow \mathbb{R}$ as follows:

$$\underline{P}_t(w) = L_t(w) + \mu [V_t(w_t) + \nabla V_t(w_t)^\top (w - w_t)]. \quad (54)$$

From [45] it is known that V_t is convex. (For the paper to be self-contained, we reprove this result in Lemma C.4.) As such, $\underline{P}_t(w) \leq P_t(w)$ for all w . Since $\underline{P}_t(w_t) = P_t(w_t)$, we obtain the estimates

$$\text{Miss}_t \leq \underline{P}_t(w_t) - \underline{P}_t(w_{t+1}) \leq \underline{P}_t(w_t) - \min_{w \in \Delta_d} \underline{P}_t(w). \quad (55)$$

Next, let $\underline{P}_t^1 : \mathbb{R}^{d-1} \rightarrow \overline{\mathbb{R}}$ be the affine reparametrization (cf. (10)) of \underline{P}_t restricted to \mathbb{A}_d , namely:

$$\underline{P}_t^1(v) := \underline{P}_t(Av + e_d). \quad (56)$$

Clearly, \underline{P}_t^1 has $\text{int}(\Delta_{d-1})$, cf. (7), as its domain, and $v \mapsto Av + e_d$ is a bijection from Δ_d to Δ_{d-1} . As a result, $\min_{w \in \Delta_d} \underline{P}_t(w) = \min_{v \in \Delta_{d-1}} \underline{P}_t^1(v)$; recalling (55), this results in

$$\text{Miss}_t \leq \underline{P}_t^1(v_t) - \min_{v \in \Delta_{d-1}} \underline{P}_t^1(v). \quad (57)$$

Function $\underline{P}_t^1(\cdot)$ is standard self-concordant in the sense of [37, Section 4.1], or 1-self-concordant in terms of our Definition 1 in Appendix A.3. (The reader may consult Appendix A.3 for the essential results on self-concordance.) As such, upper-bounding its Newton decrement

$$\underline{\text{Decr}}_t := \left\| \nabla \underline{P}_t^1(v_t) \right\|_{\nabla^2 \underline{P}_t^1(v_t)^{-1}} \quad (58)$$

with some $c < 1$ would suffice for showing that $\underline{P}_t^1(v_t) - \min_{v \in \Delta_{d-1}} \underline{P}_t^1(v) \leq C \underline{\text{Decr}}_t^2$, where C depends only on c ; in particular, Lemma A.7 shows that $c = 0.5$ allows for $C = 0.8$. Therefore, it remains to prove that $\underline{\text{Decr}}_t \leq 0.5$, while lower-bounding the right-hand side of (53) with $0.8 \underline{\text{Decr}}_t^2$.

4^o. Since $\nabla^2 \underline{P}_t^1(v) = A^\top \nabla^2 \underline{P}_t(Av + e_d)A$, cf. (11), and $\nabla^2 \underline{P}_t(w_t) = \nabla^2 L_t(w_t)$, cf. (54), one has

$$\nabla^2 \underline{P}_t^1(v_t) = \hat{\mathbf{H}}_t. \quad (59)$$

Meanwhile, since $v_t = \operatorname{argmin}_{v \in \Delta_{d-1}} P_{t-1}^{\mathbb{1}}(v)$ belongs to $\operatorname{int}(\Delta_{d-1})$, we get $\nabla P_{t-1}^{\mathbb{1}}(v_t) = 0$. Defining

$$D_t(w) := V_t(w) - V_{t-1}(w) \quad (60)$$

we thence arrive at

$$\begin{aligned} \nabla \underline{P}_t^{\mathbb{1}}(v_t) &\stackrel{(54)}{=} \nabla P_t^{\mathbb{1}}(v_t) = \nabla P_t^{\mathbb{1}}(v_t) - \nabla P_{t-1}^{\mathbb{1}}(v_t) = A^\top [\nabla P_t(w_t) - \nabla P_{t-1}(w_t)] \\ &= A^\top [\nabla L_t(w_t) - \nabla L_{t-1}(w_t) + \mu \nabla D_t(w_t)] \\ &= \hat{\nabla}_t + \mu A^\top \nabla D_t(w_t). \end{aligned} \quad (61)$$

Plugging (59) and (61) into (58) we conclude that

$$\underline{\operatorname{Decr}}_t^2 = \left\| \hat{\nabla}_t + \mu A^\top \nabla D_t(w_t) \right\|_{\hat{\mathbf{H}}_t^{-1}}^2. \quad (62)$$

Let us now focus on $\nabla D_t(w_t)$. Observe that, for any $w \in \mathbb{R}_{++}^d$,

$$-D_t(w) = \frac{1}{2} \log \det \left(\mathbf{I}_{d-1} - \nabla_t(w) \nabla_t(w)^\top \mathbf{H}_t(w)^{-1} \right) = \frac{1}{2} \log (1 - \pi_t(w)) \quad (63)$$

where we used that $\nabla_t(w) \nabla_t(w)^\top$ is rank-one. Whence by the composition formula:

$$\nabla D_t(w) = \frac{1}{2(1 - \pi_t(w))} \nabla \pi_t(w). \quad (64)$$

Finally, a rather tedious computation allows to differentiate $\pi_t(w)$ explicitly (see Lemma C.3):

$$\frac{1}{2} \nabla \pi_t(w) = \pi_t(w) \nabla \ell_t(w) - \sum_{\tau \in [t]^+} \pi_{t,\tau}(w)^2 \nabla \ell_\tau(w). \quad (65)$$

Collecting (62)–(65) we arrive at the following estimate:

$$\underline{\operatorname{Decr}}_t^2 \leq \frac{1}{(1 - \hat{\pi}_t)^2} \left\| \left[1 + (\mu - 1) \hat{\pi}_t \right] \hat{\nabla}_t - \mu \left(\sum_{\tau \in [t]^+} \hat{\pi}_{t,\tau}^2 \hat{\nabla}_\tau \right) \right\|_{\hat{\mathbf{H}}_t^{-1}}^2. \quad (66)$$

5^o. Let $\hat{\alpha}_t := 1 + (\mu - 1) \hat{\pi}_t$, then developing the square in the right-hand side of (66) results in

$$\underline{\operatorname{Decr}}_t^2 \leq \frac{\hat{\alpha}_t^2 \hat{\pi}_t + 2\mu \hat{\alpha}_t |\hat{\mathbf{E}}_1| + \mu^2 \hat{\mathbf{E}}_2}{(1 - \hat{\pi}_t)^2} \quad (67)$$

with $\hat{\mathbf{E}}_1$ and $\hat{\mathbf{E}}_2$ as follows (cf. (47)–(50)):

$$\hat{\mathbf{E}}_1 \left[:= \sum_{\tau \in [t]^+} \langle \hat{\nabla}_t, \hat{\nabla}_\tau \rangle_{\hat{\mathbf{H}}_t^{-1}} \hat{\pi}_{t,\tau}^2 \right] = \sum_{\tau \in [t]^+} \frac{1}{\sqrt{\lambda_\tau}} \hat{\pi}_{t,\tau}^3, \quad (68)$$

$$\hat{\mathbf{E}}_2 \left[:= \sum_{\tau, \nu \in [t]^+} \langle \hat{\nabla}_\tau, \hat{\nabla}_\nu \rangle_{\hat{\mathbf{H}}_t^{-1}} \hat{\pi}_{t,\tau}^2 \hat{\pi}_{t,\nu}^2 \right] = \sum_{\tau, \nu \in [t]^+} \frac{\hat{\pi}_{\tau,\nu}}{\sqrt{\lambda_\tau \lambda_\nu}} \hat{\pi}_{t,\tau}^2 \hat{\pi}_{t,\nu}^2. \quad (69)$$

Now, by Lemma C.1 proved in the appendix, we have the following bounds for the entries of $\hat{\Pi}$:

$$|\hat{\pi}_{t,\tau}| \leq \frac{1}{1+\lambda} \quad \forall \tau \in [t]; \quad |\hat{\pi}_{t,\tau}| \leq \frac{1}{\sqrt{1+\lambda}} \quad \forall \tau < 0; \quad \frac{|\hat{\pi}_{\tau,\nu}|}{\sqrt{\lambda_\tau \lambda_\nu}} \leq \frac{1}{\lambda} \quad \forall \tau, \nu \in [t]^+. \quad (70)$$

The first bound implies that $\hat{\alpha}_t \leq \frac{\lambda+\mu}{1+\lambda}$. Combining the first two bounds in (70) with (51) we get

$$|\hat{E}_1| \leq \frac{1}{\lambda} \sum_{\tau \in [t]^+} \hat{\pi}_{t,\tau}^2 \stackrel{(51)}{=} \frac{\hat{\pi}_t}{\lambda}. \quad (71)$$

Finally, applying the third and then the first bound in (70), we have that

$$\hat{E}_2 \leq \frac{1}{\lambda} \left(\sum_{\tau \in [t]^+} \hat{\pi}_{t,\tau}^2 \right)^2 \stackrel{(51)}{=} \frac{\hat{\pi}_t^2}{\lambda} \leq \frac{\hat{\pi}_t}{\lambda(1+\lambda)}. \quad (72)$$

Plugging (71)–(72) and the bound for $\hat{\alpha}_t$ back into (67) leads to

$$\underline{\text{Decr}}_t^2 \leq \left(\frac{1+\lambda}{\lambda} \right)^2 \left[\left(\frac{\lambda+\mu}{1+\lambda} \right)^2 + \frac{2\mu(\lambda+\mu)}{\lambda(1+\lambda)} + \frac{\mu^2}{\lambda(1+\lambda)} \right] \hat{\pi}_t \leq \frac{(1+\lambda)(\lambda+2\mu)^2}{\lambda^3} \hat{\pi}_t. \quad (73)$$

From the first bound in (70) we get $\underline{\text{Decr}}_t^2 \leq \frac{1}{\lambda}(1 + \frac{2\mu}{\lambda})^2$, thus $\underline{\text{Decr}}_t^2 \leq \min\{\frac{1}{4}, \frac{5\mu}{8}\hat{\pi}_t\}$ under the premise of the theorem. By Lemma A.7 and (57), this gives $\text{Miss}_t \leq 0.8 \underline{\text{Decr}}_t^2 \leq \frac{\mu}{2} \hat{\pi}_t$, cf. (53), leading to (37). The theorem is proved. \square

4 Efficient implementation

We are now about to present an implementation of (VB-FTRL) based on a quasi-Newton method run from the previous portfolio. Let us first explain why this choice is natural in our context.

First, as we verify in Appendix C, function P_t^1 , to be minimized in each round of VB-FTRL after the reparametrization (10), is self-concordant. Such functions were initially proposed in [38] as the ones for which Newton’s method has a global convergence guarantee. Moreover, it converges quadratically when initialized at a point with a small Newton decrement. On the other hand, when proving Theorem 2.1 we have shown (when bounding Decr_t) that portfolio w_t gives a small Newton decrement in terms of P_{t-1}^1 . As such, Newton’s method run from $\tilde{w}_t \approx w_t [= \text{argmin}_{w \in \Delta_d} P_{t-1}(w)]$ will produce an accurate approximation \tilde{w}_{t+1} of w_{t+1} in just a few steps, enforcing that $\tilde{w}_{t+1} \approx w_{t+1}$.

Furthermore, the convergence guarantees for Newton’s method, when run on a self-concordant function, are *affine-invariant*—just as the regret bound established in Theorem 2.1. This would not be so if we used, instead, a first-order method such as gradient descent. Given a substantial research effort that had been spent in obtaining affine-invariant regret guarantees for various algorithms,⁶ it would be quite unsatisfactory if such a regret bound—even an optimal one—came only at the price of a computational guarantee that lacks affine-invariance. (We stress this point, as it seems to be rather poorly understood in the theoretical machine learning and online optimization communities.)

⁶Specifically, getting rid of the gradient norm parameters G_2, G_∞ ; cf. Table 1 and the accompanying discussion.

Finally, we have at our disposal a computationally cheap approximate second-order oracle for P_t (and thus for $P_t^{\mathbb{1}}$ as well). Namely, it can be verified (see Lemma C.4) that the gradient of $V_t(w)$ is

$$\nabla V_t(w) = \sum_{\tau \in [t]^+} \pi_\tau(w) \nabla \ell_\tau(w), \quad (74)$$

and the Hessian of V_t satisfies the bounds $\mathbf{Q}_t(w) \preceq \nabla^2 V_t(w) \preceq 3\mathbf{Q}_t(w)$ with $\mathbf{Q}_t(w)$ as follows:

$$\mathbf{Q}_t(w) = \sum_{\tau \in [t]^+} \pi_\tau(w) \nabla^2 \ell_\tau(w). \quad (75)$$

Now, the point here is that $\mathbf{Q}_t(w)$ can be computed in $O(d^2(T+d))$ arithmetic operations (a.o.'s), whereas computing $\nabla^2 V_t(w)$ exactly would take $O(d^2(T+d)^2)$ a.o.'s. (The latter can be seen from the explicit representation of $\nabla^2 V_t$ as the sum of $(T+d)^2$ outer-product terms, cf. (121) in Appendix C.) On the other hand, it can be shown that a quasi-Newton method in which the exact Hessian of a self-concordant objective is replaced with such an approximation—i.e., the one with a constant relative error as is the case for $\nabla^2 V_t(w)$ and $\mathbf{Q}_t(w)$ —admits strong convergence guarantees—namely, local linear convergence in a Dikin ellipsoid of the optimum; see Lemma A.10.

Conceptual procedure. The above discussion suggests to replace (VB-FTRL) with the update shown in Fig. 4, and referred to as (VB-FTRL-qN) later on. In (VB-FTRL-qN), the next portfolio \tilde{w}_t is obtained by running S steps of a quasi-Newton method for the minimization problem in (VB-FTRL), under the affine reparametrization of Δ_d onto Δ_{d-1} , cf. (10)–(12), starting from the previously chosen portfolio \tilde{w}_{t-1} , and approximating the Hessian $\nabla^2 P_t(w)$ with $\nabla^2 L_t(w) + 3\mu\mathbf{Q}_t(w)$.

Input: $\tilde{w}_{t-1} \in \Delta_d$.

1. Initialize $w_t^{(0)} = \tilde{w}_{t-1}$.
2. For $s = 0, 1, \dots, S-1$ do:

(VB-FTRL-qN)

$$\begin{aligned} \mathbf{M}_{t-1}^{(s)} &= A^\top [\nabla^2 L_t(w_t^{(s)}) + 3\mu\mathbf{Q}_t(w_t^{(s)})] A, \\ w_t^{(s+1)} &= w_t^{(s)} - A[\mathbf{M}_{t-1}^{(s)}]^{-1} A^\top \nabla P_t(w_t^{(s)}). \end{aligned}$$

Output: $\tilde{w}_t = w_t^{(S)}$.

Figure 1: Conceptual implementation of (VB-FTRL) via a quasi-Newton method.

Computational guarantees. Algorithm 1 is a concrete implementation of the update rule in (VB-FTRL-qN), with $O(dT + d^2)$ overall memory use, and $O(d^2T + d^3)$ runtime per one Newton step. As we are about to see next (in Theorem 4.1), it suffices to perform $O(\log(T+d))$ Newton steps to match the regret guarantee of Theorem 2.1 up to a constant factor; thus, the overall runtime is $\tilde{O}(d^2T + d^3)$ per round.⁷ The runtime and memory use estimates we have just

⁷More precisely, round t takes $O(d^2t + d^3)$ a.o.'s, but this gives $O(d^2T + d^3)$ per round *on average over T rounds*.

announced can be verified by a line-by-line inspection of Algorithm 1. Indeed, in each round t we start by translating w_{t-1} into $v_{t-1} = v_{t-1}^{(0)}$ as per (12) to translate the problem on Δ_{d-1} a set with nonempty interior; clearly, this can be done in $O(d)$. Next, when given the previous iterate $v_t^{(s)}$, we compute the individual loss gradients $\nabla_\tau = \nabla \ell_\tau^1(v_t^{(s)})$ for each datapoint $\tau \in [t-1]^+$ (including the “fictional” datapoints $\{e_1, \dots, e_d\}$ corresponding to the logarithmic-barrier term); this takes $O(dt + d^2)$ a.o.’s in total, and requires $O(dt + d^2)$ of storage space. These data are then used to compute the Hessian $\mathbf{H}_{t-1}^{(s)} = \nabla^2 L_{t-1}^1(v_t^{(s)})$ as the sum of $O(t + d)$ one-rank matrices of the form $\lambda^{1\{\tau < 0\}} \nabla_\tau^{(s)} [\nabla_\tau^{(s)}]^\top$, in $O(td^2 + d^3)$ a.o.’s and $O(d^2)$ additional storage. After that, we compute the inverse of $\mathbf{H}_{t-1}^{(s)}$ with $O(d^3)$ runtime and $O(d^2)$ additional storage, and use $[\mathbf{H}_{t-1}^{(s)}]^{-1}$ to compute $O(t + d)$ leverage scores $\pi_\tau^{(s)}$, cf. line 11 of the algorithm, in $O(d^2)$ a.o.’s per each $\pi_\tau^{(s)}$. Using these, we then compute $\nabla P_{t-1}^1(v_t^{(s)})$, compute and invert $\mathbf{M}_{t-1}^{(s)}$ (cf. line 15) that satisfies

$$\frac{1}{3} \mathbf{M}_{t-1}^{(s)} \preceq \nabla^2 P_{t-1}^1(v_t^{(s)}) \preceq \mathbf{M}_{t-1}^{(s)}, \quad (76)$$

and perform a matrix-vector multiplication (line 16). This takes $O(td^2 + d^3)$ time and $O(d^2)$ space.

4.1 Regret guarantee for VB-FTRL-qN

As we are about to see, performing a logarithmic in $T + d$ number of quasi-Newton steps, instead of full minimization as per (VB-FTRL), suffices to match the regret bound of Theorem 2.1. More precisely, we have the following result for the regret of VB-FTRL-qN (a.k.a. Algorithm 1).

Theorem 4.1. *For any $T \in \mathbb{N}$ and market realization $x_{1:T}$, Algorithm 1 produces $\tilde{w}_{1:T}$ such that*

$$\mathcal{R}_T(\tilde{w}_{1:T} | x_{1:T}) \leq (\lambda + 2\mu) [(d-1) \log(T + \lambda d) + 1]$$

if we choose λ, μ, S satisfying the following conditions: $\lambda \geq 2e$,

$$\frac{1}{\lambda} \left(1 + \frac{2\mu}{\lambda}\right)^2 \leq \min \left\{ \frac{1}{556}, \frac{5\mu}{8(1+\lambda)} \right\}, \quad S \geq 9 \log \left(2 \max \left\{ (T + d + 1)^2, 10^4 \sqrt{1 + 3\mu} \right\} \right). \quad (77)$$

In particular, choosing $\lambda = 560$, $\mu = 2$, and $S = 18 \lceil \log(T + d + 164) \rceil + 7$ allows to guarantee that

$$\mathcal{R}_T(\tilde{w}_{1:T} | x_{1:T}) \leq 564d \log(T + 560d).$$

Proof. As previously noted, Algorithm 1 precisely corresponds to the updates in (VB-FTRL-qN). Let $\tilde{w}_{1:T}$ and $w_{1:T}$ be two sequences produced by (VB-FTRL-qN) and (VB-FTRL) correspondingly on the same $x_{1:T}$. Let also $w_0 = \frac{1}{d} \mathbf{1}_d$, and observe that $\tilde{w}_0 = w_0 = w_1$ is the exact minimizer of P_0 . The following technical result claims, in a nutshell, that \tilde{w}_t remains very close to w_t for all $t \in [T]$.

Lemma 4.1. *Let λ, μ, S be as in the premise of Theorem 4.1. Then for all $t \in \{0\} \cup [T]$, one has*

$$\|\tilde{w}_t - w_t\|_{\nabla^2 L_t(w_t)} \leq \min \left\{ \frac{1}{(T + d + 1)^2}, 10^{-4} (1 + 3\mu)^{-1/2} \right\}. \quad (78)$$

Algorithm 1 VB-FTRL-qN : Efficient implementation of VB-FTRL via a quasi-Newton method

Input: $\lambda > 0, \mu > 0, S \in \mathbb{N}$

```

1:  $\tilde{w}_0 := \frac{1}{d} \mathbf{1}_d$   $\triangleright \tilde{w}_0 = \operatorname{argmin}_{w \in \Delta_d} P_0(w)$ 
2: for  $t \in [T]$  do
3:    $v_t^{(s)} := A^+(\tilde{w}_{t-1} - e_d)$   $\triangleright \operatorname{map} \tilde{w}_{t-1} \in \Delta_d$  to  $\Delta_{d-1}$ , cf. (12)
4:   for  $s \in \{0, \dots, S-1\}$  do
5:     for  $\tau \in [t-1]^+$  do  $\triangleright [t-1]^+ = [t-1] \cup \{-1, \dots, -d\}$ 
6:        $\lambda_\tau := \lambda \mathbf{1}_{\{\tau < 0\}}$   $\triangleright \text{recall that } x_{-j} = e_j \text{ for } j \in [d]$ 
7:        $\nabla_\tau^{(s)} := -\frac{1}{x_\tau^\top (Av_t^{(s)} + e_d)} A^\top x_\tau$   $\triangleright \nabla_\tau^{(s)} = \nabla \ell_\tau^{\mathbf{1}}(v_t^{(s)})$ 
8:     end for
9:      $\mathbf{H}_{t-1}^{(s)} := \sum_{\tau \in [t-1]^+} \lambda_\tau \nabla_\tau^{(s)} [\nabla_\tau^{(s)}]^\top$   $\triangleright \mathbf{H}_{t-1}^{(s)} = \nabla^2 L_{t-1}^{\mathbf{1}}(v_t^{(s)})$ 
10:     $\Sigma_{t-1}^{(s)} := [\mathbf{H}_{t-1}^{(s)}]^{-1}$ 
11:    for  $\tau \in [t-1]^+$  do
12:       $\pi_\tau^{(s)} := \lambda_\tau \left\| \nabla_\tau^{(s)} \right\|_{\Sigma_{t-1}^{(s)}}^2$ 
13:    end for
14:     $\mathbf{g}_{t-1}^{(s)} := \sum_{\tau \in [t-1]^+} (\lambda_\tau + \mu \pi_\tau^{(s)}) \nabla_\tau^{(s)}$   $\triangleright \mathbf{g}_{t-1}^{(s)} = \nabla P_{t-1}^{\mathbf{1}}(v_t^{(s)})$ 
15:     $\mathbf{M}_{t-1}^{(s)} := \mathbf{H}_{t-1}^{(s)} + 3\mu \sum_{\tau \in [t-1]^+} \pi_\tau^{(s)} \nabla_\tau^{(s)} [\nabla_\tau^{(s)}]^\top$ 
16:     $v_t^{(s+1)} := v_t^{(s)} - [\mathbf{M}_{t-1}^{(s)}]^{-1} \mathbf{g}_{t-1}^{(s)}$   $\triangleright \text{quasi-Newton step for } P_{t-1}^{\mathbf{1}}(\cdot)$ 
17:  end for
18:  Predict  $\tilde{w}_t := Av_t^{(S)} + e_d$   $\triangleright \operatorname{map} v_t^{(S)} \in \Delta_{d-1}$  back to  $\Delta_d$ 
19:  Receive  $x_t \in \mathbb{R}_+^d$ 
20: end for

```

Deferring the proof of Lemma 4.1 to Appendix D we shall now focus on proving the theorem.

1^o. Mimicking (35) we get

$$\begin{aligned}
\mathcal{R}_T(\tilde{w}_{1:T} | x_{1:T}) &\leq \operatorname{Bias}_T + \sum_{t \in [T]} \left[\ell_t(\tilde{w}_t) + \min_{w \in \Delta_d} P_{t-1}(w) - \min_{w \in \Delta_d} P_t(w) \right] \\
&\leq \operatorname{Bias}_T + \sum_{t \in [T]} \left[\ell_t(\tilde{w}_t) + P_{t-1}(\tilde{w}_t) - P_t(w_{t+1}) \right] \\
&\leq \operatorname{Bias}_T + \sum_{t \in [T]} [\widetilde{\operatorname{Miss}}_t + \widetilde{\operatorname{Gain}}_t]
\end{aligned} \tag{79}$$

where the term Bias_T is the same as in the proof of Theorem 2.1, and the terms summed over t are

$$\widetilde{\operatorname{Miss}}_t := P_t(\tilde{w}_t) - P_t(w_{t+1}), \quad \widetilde{\operatorname{Gain}}_t := \mu [V_{t-1}(\tilde{w}_t) - V_t(\tilde{w}_t)]. \tag{80}$$

Since $\lambda \geq 2e$, Bias_T admits the same bound (36) as before. Moreover, for $\operatorname{Miss}_t = P_t(w_t) - P_t(w_{t+1})$ and $\operatorname{Gain}_t = \mu [V_{t-1}(w_t) - V_t(w_t)]$ we still have that $\operatorname{Miss}_t + \operatorname{Gain}_t \leq 0$ for all $t \in [T]$,

simply because these terms, just like (36), depend only on $w_{1:T}$ (and not $\tilde{w}_{1:T}$), and the constraint on λ, μ imposed in (77) is stronger than its counterpart in Theorem 2.1. Thus, it remains to show that $\sum_{t \in [T]} \widetilde{\text{Error}}_t \leq \lambda + 2\mu$ for the sum of approximation errors $\widetilde{\text{Error}}_t := \widetilde{\text{Miss}}_t - \text{Miss}_t + \widetilde{\text{Gain}}_t - \text{Gain}_t$.
 2° . Clearly,

$$\widetilde{\text{Error}}_t = L_t(\tilde{w}_t) - L_t(w_t) + \mu[V_{t-1}(\tilde{w}_t) - V_{t-1}(w_t)].$$

By Lemma 4.1, Lemma A.9, and since $L_t^\sharp(\cdot)$ and $L_{t-1}^\sharp(\cdot)$ are 1-self-concordant, cf. Corollary A.1,

$$(1 - \varepsilon)^2 \mathbf{H}_t(w_t) \preceq \mathbf{H}_t(\tilde{w}_t) \preceq \frac{1}{(1 - \varepsilon)^2} \mathbf{H}_t(w_t),$$

$$(1 - \varepsilon)^2 \mathbf{H}_{t-1}(w_t) \preceq \mathbf{H}_{t-1}(\tilde{w}_t) \preceq \frac{1}{(1 - \varepsilon)^2} \mathbf{H}_{t-1}(w_t)$$

where ε is the right-hand side of (78), and the second line is due to $\|\tilde{w}_t - w_t\|_{\nabla^2 L_{t-1}(w_t)} \leq \varepsilon$. Thus,

$$\begin{aligned} \mu[V_{t-1}(\tilde{w}_t) - V_{t-1}(w_t)] &= \frac{\mu}{2} \log \det \left(\mathbf{H}_{t-1}(w_t)^{-1/2} \mathbf{H}_{t-1}(\tilde{w}_t) \mathbf{H}_{t-1}(w_t)^{-1/2} \right) \leq \mu(d-1) \log \left(\frac{1}{1 - \varepsilon} \right) \\ &\leq 2\mu(d-1)\varepsilon. \end{aligned}$$

where in the end we used that $\frac{1}{1-u} \leq 1 + 2u$ for $u \in [0, \frac{1}{2}]$. On the other hand, applying Lemma A.8 to $L_t^\sharp(\cdot)$ and noting that $\|\nabla L_t(w)\|_{\nabla^2 L_t(w)^{-1}} \leq t + d$ by triangle inequality, we get

$$L_t(\tilde{w}_t) - L_t(w_t) \leq \nabla L_t(w_t)^\top (\tilde{w}_t - w_t) + 0.8 \|\tilde{w}_t - w_t\|_{\nabla^2 L_t(w_t)}^2 \leq (T + d)\varepsilon + 0.8\varepsilon^2 \leq (T + d + 1)\varepsilon.$$

As such, we have a (very conservative) bound $\sum_{t \in [T]} \widetilde{\text{Error}}_t \leq (\lambda + 2\mu)(T + d + 1)^2 \varepsilon \leq \lambda + 2\mu$. \square

Acknowledgments

Dmitrii M. Ostrovskii acknowledges support by the National Science Foundation grant CIF-1908905.

A Auxiliary results

A.1 Logarithmic barrier and smoothed portfolios

Lemma A.1. *Let $w_T^* = \operatorname{argmin}_{w \in \Delta_d} L_T(w)$, then*

$$w_T^*[i] \geq \frac{\lambda}{T + \lambda d}, \quad \forall i \in [d]. \quad (81)$$

Proof. Let $1/w$ be the vector with entries $w[i] = 1/w[i]$. The first-order optimality conditions read

$$(w - w_T^*)^\top \left(\frac{\lambda}{w_T^*} + \sum_{t \in [T]} \frac{x_t}{x_t^\top w_T^*} \right) \leq 0, \quad \forall w \in \Delta_d.$$

Rearranging, we conclude that

$$\lambda \sum_{i \in [d]} \frac{w[i]}{w_T^*[i]} + \sum_{t \in [T]} \frac{x_t^\top w}{x_t^\top w_T^*} \leq \lambda d + T.$$

The second sum is nonnegative; choosing $w = e_i$ for $i \in [d]$ violating (81) gives a contradiction. \square

Lemma A.2. For all $w \in \Delta_d$, the smoothed portfolio $[w]_\alpha := (1 - \alpha)w + \frac{\alpha}{d}\mathbf{1}_d$, $\alpha \in [0, 1)$, satisfies

$$\sum_{t \in [T]} \ell_t([w]_\alpha) - \ell_t(w) \leq \frac{\alpha}{1 - \alpha} T.$$

In particular, for $\alpha = \frac{\lambda(d-1)}{T + \lambda(d-1)}$ we have that

$$\sum_{t \in [T]} \ell_t([w]_\alpha) - \ell_t(w) \leq \lambda(d-1).$$

Proof. Observe that $w - [w]_\alpha = \frac{\alpha}{1-\alpha}([w]_\alpha - \frac{1}{d}\mathbf{1}_d)$. Whence by convexity we get

$$\sum_{t \in [T]} \ell_t([w]_\alpha) - \ell_t(w) \leq \sum_{t \in [T]} \frac{x_t^\top (w - [w]_\alpha)}{x_t^\top [w]_\alpha} = \frac{\alpha}{1 - \alpha} \sum_{t \in [T]} \left(1 - \frac{1}{d} \cdot \frac{\mathbf{1}_d^\top x_t}{x_t^\top [w]_\alpha}\right) \leq \frac{\alpha}{1 - \alpha} T. \quad \square$$

Lemma A.3. For all $w \in \mathbb{R}_{++}^d$ and $x \in \mathbb{R}_+^d$, function $R(\cdot)$ in (14) satisfies $\nabla^2 R(w) \succcurlyeq \frac{1}{(x^\top w)^2} x x^\top$.

Proof. Equivalently, we must prove that, for w, x as in the premise of the lemma, and for any $u \in \mathbb{R}^d$,

$$(x^\top w)^2 \sum_{i \in [d]} \frac{u[i]^2}{w[i]^2} \geq (x^\top u)^2.$$

Clearly, it suffices to consider $u \in \mathbb{R}_+^d$. Let $r \in \mathbb{R}_+^d$ have entries $r[i] = \frac{u[i]}{w[i]}$, then we have that

$$\sqrt{\sum_{i \in [d]} \frac{u[i]^2}{w[i]^2}} x^\top w = \|r\|_2 x^\top w \geq \|r\|_\infty x^\top w \geq \sum_{i \in [d]} r[i] x[i] w[i] = x^\top u. \quad \square$$

A.2 Facts from linear algebra

Lemma A.4 ([20, Lemma 12]). For any $\mathbf{H} \succ \mathbf{H}' \succcurlyeq 0$, it holds that $\log \left(\frac{\det(\mathbf{H})}{\det(\mathbf{H} - \mathbf{H}')} \right) \geq \text{tr}(\mathbf{H}^{-1} \mathbf{H}')$.

In particular, if \mathbf{H}' is rank-one, i.e. $\mathbf{H} = uu^\top$ for some vector u , then $\log \left(\frac{\det(\mathbf{H})}{\det(\mathbf{H} - uu^\top)} \right) \geq \|u\|_{\mathbf{H}^{-1}}^2$.

Proof. The result follows by using that $\log \det(\cdot)$ is concave on the positive-semidefinite cone (see [8]):

$$\log \left(\frac{\det(\mathbf{H})}{\det(\mathbf{H} - \mathbf{H}')} \right) = \log \det(\mathbf{H}(\mathbf{H} - \mathbf{H}')^{-1}) \geq \text{tr}(\mathbf{I} - (\mathbf{H} - \mathbf{H}')\mathbf{H}^{-1}) = \text{tr}(\mathbf{H}^{-1} \mathbf{H}'). \quad \square$$

Lemma A.5. Let $\mathbf{H} \in \mathbb{S}_{++}^n$ and $u, v \in \mathbb{R}^n$. Then $\nabla_{\mathbf{H}}(\langle u, v \rangle_{\mathbf{H}^{-1}}) = -\frac{1}{2}\mathbf{H}^{-1}(uv^\top + vu^\top)\mathbf{H}^{-1}$. In particular,

$$\nabla_{\mathbf{H}}(\|u\|_{\mathbf{H}^{-1}}^2) = -\mathbf{H}^{-1}uu^\top\mathbf{H}^{-1}.$$

Proof. Fix an arbitrary symmetric matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, and consider $t \in \mathbb{R}$ for which $\mathbf{H} - t\mathbf{M} \succ 0$. In terms of the Cholesky factorization $\mathbf{H} = \mathbf{R}\mathbf{R}^\top$ we have that, as $t \rightarrow 0$,

$$(\mathbf{H} + t\mathbf{M})^{-1} = \mathbf{R}^{-\top}(\mathbf{I} - t\mathbf{R}^{-1}\mathbf{M}\mathbf{R}^{-\top})\mathbf{R}^{-1} + o(t) = \mathbf{H}^{-1} - t\mathbf{H}^{-1}\mathbf{M}\mathbf{H}^{-1} + o(t).$$

Whence $u^\top \mathbf{H}^{-1}u - u^\top (\mathbf{H} + t\mathbf{M})^{-1}u = t \text{tr}(\mathbf{M}\mathbf{H}^{-1}uu^\top \mathbf{H}^{-1}) + o(t)$, and the formula for $\nabla_{\mathbf{H}}(\|u\|_{\mathbf{H}^{-1}}^2)$ is verified. More generally,

$$u^\top \mathbf{H}^{-1}v - u^\top (\mathbf{H} + t\mathbf{M})^{-1}v = t \text{tr}(\mathbf{M}\mathbf{H}^{-1}vu^\top \mathbf{H}^{-1}) + o(t) = \frac{t}{2} \text{tr}(\mathbf{M}\mathbf{H}^{-1}(uv^\top + vu^\top)\mathbf{H}^{-1}) + o(t).$$

Since $\mathbf{H}^{-1}(uv^\top + vu^\top)\mathbf{H}^{-1}$ is a symmetric matrix, this verifies the formula for $\nabla_{\mathbf{H}}(\langle u, v \rangle_{\mathbf{H}^{-1}})$. \square

A.3 Self-concordant functions

In this section, we recall the notion of self-concordant functions and the key properties of such functions. This class of functions was thoroughly studied in the seminal work of Nesterov and Nemirovski [38] in connection with interior-point methods; earlier, Vaidya used similar tools in the context of cutting-plane methods [45]. The results presented next are well known; their proofs can be found in [37, Section 4], see also [40, Proposition B.2] for a more concise proof of Lemma A.9.

Definition 1. *Let $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ be convex, proper, and three times differentiable on its domain $\text{dom}(f) := \{z : f(z) < \infty\}$ which is open, has nonempty interior, and does not contain straight lines. Then f is called M -self-concordant if it is a barrier on its domain—i.e. $f(z) \rightarrow_{z \rightarrow z_0} \infty$ for any z_0 in the boundary of $\text{dom}(f)$ —and the following inequality holds for all $z \in \text{dom}f$ and $u \in \mathbb{R}^n$:*

$$|\nabla^3 f(z)[u, u, u]| \leq 2M (\nabla^2 f(z)[u, u])^{3/2}. \quad (82)$$

The following property of self-concordant functions is well-known, see e.g. [37, Section 4.1].

Lemma A.6 ([37, Theorem 4.1.1]). *Let f_1 and f_2 be M_1 - and M_2 -self-concordant, respectively. For $\lambda_1, \lambda_2 > 0$, the function $\lambda_1 f_1 + \lambda_2 f_2$ has domain $\text{dom}(f_1) \cap \text{dom}(f_2)$ and is M -self-concordant,*

$$M = \max \left\{ \frac{M_1}{\sqrt{\lambda_1}}, \frac{M_2}{\sqrt{\lambda_2}} \right\}.$$

Lemma A.6 implies self-concordance of the reparametrized logarithmic-barrier function $L_t^\mathbb{1}(\cdot)$.

Corollary A.1. *Assume $\lambda \geq 1$, then the restriction $L_t^\mathbb{1} : \mathbb{R}^{d-1} \rightarrow \overline{\mathbb{R}}$ of L_t to \mathbb{A}_d as in (9), namely*

$$L_t^\mathbb{1}(v) = L_t(Av + e_d)$$

with A defined in (8), is 1-self-concordant according to Definition 1 with domain $\text{int}(\Delta_{d-1})$, cf. (7).

Proof. It is evident that $\text{dom}(L_t^\mathbb{1}) = \text{int}(\Delta_{d-1})$ where Δ_{d-1} is the “solid” simplex in \mathbb{R}^{d-1} , cf. (7). In particular, $\text{dom}(L_t^\mathbb{1})$ is open and does not contain straight lines, and $L_t^\mathbb{1} \rightarrow +\infty$ on its boundary. Thus, the requirements on the domain are verified. On the other hand, inequality (82) holds with $M = 1$ for each of the functions $\ell_\tau(Av + e_d)$, $\tau \in [d]$, and with $M = 1/\sqrt{\lambda} \leq 1$ for the functions $-\lambda \log(e_i^\top (Av + e_d))$, $i \in [d]$. Applying Jensen’s inequality as in the proof of Lemma A.6 (cf. [37, Theorem 4.1.1]) we verify (82) with $M = 1$ for $L_t^\mathbb{1}(\cdot)$. That is, $L_t^\mathbb{1}(\cdot)$ is 1-self-concordant. \square

Following [37, Section 4.1] we now define a pair of mutually conjugate functions ω, ψ on \mathbb{R}^+ by

$$\omega(r) := r - \log(1 + r), \quad \psi(r) := -r - \log(1 - r). \quad (83)$$

Note that ω and ψ are increasing, $\omega(r) \leq \frac{1}{2}r^2 \leq \psi(r)$, and these bounds are tight at $r = 0$. In fact, $\omega(r)$ is a Huber-type function, so $\omega(r) = (1 - o(1))r$ as $r \rightarrow \infty$; meanwhile, $\psi(r) \rightarrow +\infty$ as $r \uparrow 1$. More specifically, it can be verified that $\omega(r) \geq 0.3r^2$ for $r \leq 1$, and $\psi(r) \leq 0.8r^2$ for $r \leq 0.5$.

Lemma A.7 ([37, Theorems 4.1.10–4.1.11, 4.1.13]). *Let f be 1-self-concordant, and some $\hat{z} \in \text{dom}(f)$ satisfies $\|\nabla f(\hat{z})\|_{\nabla^2 f(\hat{z})^{-1}} < 1$. Then f has a unique minimizer z^* , and*

$$\omega(\|\nabla f(z)\|_{\nabla^2 f(z)^{-1}}) \leq f(z) - f(z^*) \leq \psi(\|\nabla f(z)\|_{\nabla^2 f(z)^{-1}})$$

where the lower bound holds for all $z \in \text{dom}(f)$, and the upper bound requires $\|\nabla f(z)\|_{\nabla^2 f(z)^{-1}} < 1$. In particular, $0.3 \|\nabla f(z)\|_{\nabla^2 f(z)^{-1}}^2 \leq f(z) - f(z^) \leq 0.8 \|\nabla f(z)\|_{\nabla^2 f(z)^{-1}}^2$ if $\|\nabla f(z)\|_{\nabla^2 f(z)^{-1}} \leq 0.5$.*

Lemma A.8 ([37, Theorems 4.1.5, 4.1.7–4.1.8]). *Let f be 1-self-concordant, $z \in \text{dom}(f)$, then*

$$\omega(\|z' - z\|_{\nabla^2 f(z)}) \leq f(z') - f(z) - \nabla f(z)^\top (z' - z) \leq \psi(\|z' - z\|_{\nabla^2 f(z)}).$$

Here, the lower bound holds $\forall z' \in \text{dom}(f)$ and the upper bound on 1-Dikin ellipsoid $\mathcal{E}_{f,1}(z)$ where

$$\mathcal{E}_{f,r}(z) := \{z' \in \mathbb{R}^n : \|z' - z\|_{\nabla^2 f(z)} < r\}. \quad (84)$$

In particular, $\mathcal{E}_{f,1}(z)$ is contained in $\text{dom}(f)$. Moreover, for $z' \in \mathcal{E}_{f,\frac{1}{2}}(z)$ we have the following:

$$0.3 \|z' - z\|_{\nabla^2 f(z)}^2 \leq f(z') - f(z) - \nabla f(z)^\top (z' - z) \leq 0.8 \|z' - z\|_{\nabla^2 f(z)}^2.$$

Lemma A.9 ([37, Theorem 4.1.6]). *Let f be 1-self-concordant and $z \in \text{dom}(f)$, then we have*

$$(1 - \|z' - z\|_{\nabla^2 f(z)})^2 \nabla^2 f(z) \preceq \nabla^2 f(z') \preceq \frac{1}{(1 - \|z' - z\|_{\nabla^2 f(z)})^2} \nabla^2 f(z)$$

for all $z' \in \mathcal{E}_{f,1}(z)$. In particular, for $z' \in \mathcal{E}_{f,\frac{1}{2}}(z)$ we have that $0.25 \nabla^2 f(z) \preceq \nabla^2 f(z') \preceq 4 \nabla^2 f(z)$.

A.4 Quasi-Newton method

Next we focus on a quasi-Newton method in which the actual Hessian is replaced with an approximation with a constant relative accuracy. Such a method converges linearly if initialized in a constant-radius Dikin ellipsoid of the minimizer of a self-concordant function. Analogous results were obtained in [33, Lemma 11], for a modified notion of self-concordance.

Lemma A.10 (Linear convergence for a quasi-Newton method). *Let f be 1-self-concordant, $z \in \text{dom}(f)$ be such that $\|\nabla f(z)\|_{\nabla^2 f(z)^{-1}} \leq \frac{c}{3}$ for some $c \leq 1$, and consider $z^+ := z - \tilde{\mathbf{H}}^{-1} \nabla f(z)$ with*

$$c\tilde{\mathbf{H}} \preceq \nabla^2 f(z) \preceq \tilde{\mathbf{H}}.$$

Then $z^+ \in \text{dom}(f)$, and we have that $\|\nabla f(z^+)\|_{\nabla^2 f(z^+)^{-1}} \leq (1 - \frac{c}{3}) \|\nabla f(z)\|_{\nabla^2 f(z)^{-1}}$.

Proof. Define $\mathbf{H}(z') := \nabla^2 f(z')$ for any $z' \in \text{dom}(f)$, and $\mathbf{J} := \mathbf{H}(z)^{1/2} \tilde{\mathbf{H}}^{-1} \mathbf{H}(z)^{1/2}$. Then, by the premise of the lemma, $c\mathbf{H}(z)^{-1} \preceq \tilde{\mathbf{H}}^{-1} \preceq \mathbf{H}(z)^{-1}$, which implies

$$c\mathbf{I} \preceq \mathbf{J} \preceq \mathbf{I} \quad \text{and} \quad 0 \preceq \mathbf{I} - \mathbf{J} \preceq (1 - c)\mathbf{I}.$$

As a result, defining $\delta(z') := \|\nabla f(z')\|_{\nabla^2 f(z')^{-1}}$ for brevity, we get

$$\|z^+ - z\|_{\mathbf{H}(z)}^2 = \|\tilde{\mathbf{H}}^{-1} \nabla f(z)\|_{\mathbf{H}(z)}^2 = \nabla f(z)^\top \mathbf{H}(z)^{-1/2} \mathbf{J}^2 \mathbf{H}(z)^{-1/2} \nabla f(z) \leq \delta(z)^2,$$

whence (by Lemma A.9) $z^+ \in \text{dom}(f)$ and

$$(1 - \delta(z))^2 \mathbf{H}(z)^{-1} \preceq \mathbf{H}(z^+)^{-1} \preceq \frac{1}{(1 - \delta(z))^2} \mathbf{H}(z)^{-1}. \quad (85)$$

Now, define $\mathbf{H}(z) := \nabla^2 f(z)$ and $z_s := sz^+ + (1-s)z$ for $s \in [0, 1]$. Since $\nabla f(z) + \tilde{\mathbf{H}}(z^+ - z) = 0$,

$$\begin{aligned}
\|\nabla f(z^+)\|_{\mathbf{H}(z)^{-1}} &= \left\| \nabla f(z^+) - \nabla f(z) - \tilde{\mathbf{H}}(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}} \\
&= \left\| \int_0^1 (\mathbf{H}(z_s) - \tilde{\mathbf{H}})(z^+ - z) ds \right\|_{\mathbf{H}(z)^{-1}} \\
&\leq \int_0^1 \left\| (\mathbf{H}(z_s) - \tilde{\mathbf{H}})(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}} ds \\
&\leq \int_0^1 \left\| (\mathbf{H}(z_s) - \mathbf{H}(z))(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}} ds + \left\| (\tilde{\mathbf{H}} - \mathbf{H}(z))(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}}.
\end{aligned} \tag{86}$$

For the last term in the right-hand side, we get

$$\begin{aligned}
\left\| (\tilde{\mathbf{H}} - \mathbf{H}(z))(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}}^2 &= \left\| (\mathbf{I} - \mathbf{H}(z)\tilde{\mathbf{H}}^{-1})\nabla f(z) \right\|_{\mathbf{H}(z)^{-1}}^2 \\
&= \nabla f(z)^\top \mathbf{H}(z)^{-1/2} (\mathbf{I} - \mathbf{J})^2 \mathbf{H}(z)^{-1/2} \nabla f(z) \leq (1-c)^2 \delta(z)^2.
\end{aligned} \tag{87}$$

On the other hand, since $z_s - z = s(z^+ - z)$, similarly to (85), we have that

$$(1 - s\delta(z))^2 \mathbf{H}(z)^{-1} \preceq \mathbf{H}(z_s)^{-1} \preceq \frac{1}{(1 - s\delta(z))^2} \mathbf{H}(z)^{-1},$$

and since $(1-u)^2 \leq \frac{1}{(1-u)^2}$ for $u \in (0, 1)$, we get $\left\| \mathbf{H}(z)^{1/2} \mathbf{H}(z_s)^{-1} \mathbf{H}(z)^{1/2} - \mathbf{I} \right\| \leq \frac{1}{(1-s\delta(z))^2} - 1$. As a result, since

$$\left\| (\mathbf{H}(z_s) - \mathbf{H}(z))(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}}^2 = (z^+ - z)^\top \mathbf{H}(z)^{\frac{1}{2}} (\mathbf{H}(z)^{\frac{1}{2}} \mathbf{H}(z_s)^{-1} \mathbf{H}(z)^{\frac{1}{2}} - \mathbf{I})^2 \mathbf{H}(z)^{\frac{1}{2}} (z^+ - z),$$

we get

$$\left\| (\mathbf{H}(z_s) - \mathbf{H}(z))(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}} \leq \left(\frac{1}{(1 - s\delta(z))^2} - 1 \right) \|z^+ - z\|_{\mathbf{H}(z)} \leq \left(\frac{1}{(1 - s\delta(z))^2} - 1 \right) \delta(z)$$

and estimate the integral term:

$$\int_0^1 \left\| (\mathbf{H}(z_s) - \mathbf{H}(z))(z^+ - z) \right\|_{\mathbf{H}(z)^{-1}} ds \leq \delta(z) \int_0^1 \left(\frac{1}{(1 - s\delta(z))^2} - 1 \right) ds = \frac{\delta(z)^2}{1 - \delta(z)}. \tag{88}$$

Finally, combining (85)–(88), and using that $\delta(z) \leq \frac{\epsilon}{3}$, we conclude that

$$\delta(z^+) \leq (1-c) \frac{\delta(z)}{1 - \delta(z)} + \frac{\delta(z)^2}{(1 - \delta(z))^2} \leq \left(1 - \frac{2c}{3}\right) \frac{\delta(z)}{1 - \delta(z)} \leq \frac{3-2c}{3-c} \delta(z) \leq \left(1 - \frac{c}{3}\right) \delta(z). \quad \square$$

Corollary A.2. *Let f be 1-self-concordant and minimized at z^* , and assume that $z_0 \in \text{dom}(f)$ is such that $\|\nabla f(z_0)\|_{\nabla^2 f(z_0)^{-1}} \leq \frac{\epsilon}{3}$ for some $c \leq 1$. Consider the following sequence for $s \in \{0\} \cup [S-1]$:*

$$z_{s+1} := z_s - \tilde{\mathbf{H}}_s^{-1} \nabla f(z_s) \quad \text{where} \quad c\tilde{\mathbf{H}}_s \preceq \nabla^2 f(z_s) \preceq \tilde{\mathbf{H}}_s.$$

Then, for any $\varepsilon \leq \frac{1}{4}$, we have the following inequalities if $S \geq \lceil \frac{3}{c} \log(\frac{1}{\varepsilon}) \rceil$:

$$\begin{aligned}
\max \left\{ \frac{1}{4} \|z_S - z^*\|_{\nabla^2 f(z^*)}^2, \quad f(z_S) - f(z^*), \quad \|\nabla f(z_S)\|_{\nabla^2 f(z_S)^{-1}}^2 \right\} &\leq \varepsilon^2, \\
(1 - 2\varepsilon)^2 \nabla^2 f(z^*) &\preceq \nabla^2 f(z_S) \preceq \frac{1}{(1 - 2\varepsilon)^2} \nabla^2 f(z^*).
\end{aligned}$$

Proof. Using Lemma A.10 sequentially S times, $\|\nabla f(z_S)\|_{\nabla^2 f(z_S)^{-1}} \leq \frac{\epsilon}{3} (1 - \frac{\epsilon}{3})^S \leq \min\{\epsilon, \frac{1}{3}\}$. Then Lemma A.7 gives $f(z_S) - f(z^*) \leq 0.8\epsilon^2$. Now, one can verify that $\omega(r) \geq 0.1r$ for $r \geq 0.5$, whence by Lemma A.8 first $\|z_S - z^*\|_{\nabla^2 f(z^*)} \leq \max\{0.5, 8\epsilon^2\} \leq 0.5$, and then $\|z_S - z^*\|^2 \leq \frac{8}{3}\epsilon^2$. \square

B Proofs deferred from Section 2.2

B.1 Proof of Proposition 2.1

Function $h_t(w) = e^{-\frac{1}{\mu}\ell_t(w)} [= (x_t^\top w)^{1/\mu}]$ is concave on Δ_d when $\mu \geq 1$; thus, by Jensen's inequality,

$$\begin{aligned} \ell_t(w_t) &\leq -\mu \log \mathbb{E}_{w \sim \phi_t} \left[\exp \left(-\frac{1}{\mu} \ell_t(w) \right) \right] \\ &\leq -\mu \log \left[\int_{\Delta_d} \exp \left(-\frac{1}{\mu} L_t(w) \right) dw \right] + \mu \log \left[\int_{\Delta_d} \exp \left(-\frac{1}{\mu} L_{t-1}(w) \right) dw \right]. \end{aligned}$$

Whence via telescoping and (15):

$$\sum_{t \in [T]} \ell_t(w_t) \leq -\mu \log \left[\int_{\Delta_d} \exp \left(-\frac{1}{\mu} L_T(w) \right) dw \right] + \mu \log \left[\int_{\Delta_d} \exp \left(-\frac{\lambda}{\mu} R(w) \right) dw \right]. \quad (89)$$

Now, let $w_T^* = \operatorname{argmin}_{w \in \Delta_d} L_T(w)$. Define the set $\Delta_{d,T}^* := \left\{ \frac{T}{T+1} w_T^* + \frac{1}{T+1} w, \forall w \in \Delta_d \right\}$, i.e. a copy of Δ_d shrunk by the factor of $\frac{1}{T+1}$ and shifted by w_T^* . Since $h_t(\cdot)$ is concave and positive, we have:

$$h_t(w) \geq \frac{T}{T+1} h_t(w_T^*) \quad \forall w \in \Delta_{d,T}^*.$$

As such, for all $w \in \Delta_{d,T}^*$ we have that

$$\exp \left(-\frac{1}{\mu} L_T(w) \right) \geq \exp \left(-\frac{\lambda}{\mu} R(w) \right) \cdot \left(\frac{T}{T+1} \right)^T \prod_{t \in [T]} h_t(w_T^*) \geq \frac{1}{e} \exp \left(-\frac{\lambda}{\mu} R(w) \right) \prod_{t \in [T]} h_t(w_T^*).$$

After taking the logarithm and noting that $\Delta_{d,T}^* \subseteq \Delta_d$, this results in

$$-\mu \log \left[\int_{\Delta_d} \exp \left(-\frac{1}{\mu} L_T(w) \right) dw \right] \leq -\mu \log \left[\frac{1}{e} \int_{\Delta_{d,T}^*} \exp \left(-\frac{\lambda}{\mu} R(w) \right) dw \right] + \sum_{t \in [T]} \ell_t(w_T^*).$$

Returning to (89), this results in

$$\begin{aligned} \mathcal{R}_T(w_{1:T} | x_{1:T}) &\leq \sum_{t \in [T]} \ell_t(w_T^*) - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w) \\ &\quad + \mu \log \left[\int_{\Delta_d} \exp \left(-\frac{\lambda}{\mu} R(w) \right) dw \right] - \mu \log \left[\frac{1}{e} \int_{\Delta_{d,T}^*} \exp \left(-\frac{\lambda}{\mu} R(w) \right) dw \right]. \end{aligned} \quad (90)$$

Observing that $R(w) \geq R(\frac{1}{d}\mathbf{1}_d) = d \log(d)$ on Δ_d allows to estimate the penultimate term in (90):

$$\mu \log \left[\int_{\Delta_d} \exp \left(-\frac{\lambda}{\mu} R(w) \right) dw \right] \leq \mu \log[\operatorname{Area}(\Delta_d)] - \lambda d \log d. \quad (91)$$

Meanwhile, $R(w)$, as a convex function, is maximized on a vertex of $\Delta_{d,T}^*$, whence for any $w \in \Delta_{d,T}^*$:

$$\begin{aligned} R(w) &\leq \max_{k \in [d]} R\left(\frac{T}{T+1}w_T^* + \frac{1}{T+1}e_k\right) = \max_{k \in [d]} \left\{ \log\left(\frac{T+1}{Tw_T^*[k]+1}\right) + \sum_{i \in [d] \setminus \{k\}} \log\left(\frac{T+1}{Tw_T^*[i]}\right) \right\} \\ &\leq R(w_T^*) + (d-1) \log\left(\frac{T+1}{T}\right). \end{aligned}$$

This allows to estimate the last term in (90):

$$\begin{aligned} -\mu \log \left[\frac{1}{e} \int_{\Delta_{d,T}^*} \exp\left(-\frac{\lambda}{\mu}R(w)\right) dw \right] &\leq \lambda R(w_T^*) + \lambda(d-1) \log\left(\frac{T+1}{T}\right) + \mu - \mu \log[\text{Area}(\Delta_{d,T}^*)] \\ &\leq \lambda R(w_T^*) + \lambda(d-1) \log(2) + \mu - \mu \log[\text{Area}(\Delta_d)] \\ &\quad + \mu(d-1) \log(T+1). \end{aligned} \tag{92}$$

Here we used that $\text{Area}(\Delta_{d,T}^*) = \frac{1}{(T+1)^{d-1}} \text{Area}(\Delta_d)$. Finally, proceeding as in step 1^o of the proof of Theorem 2.1 we get

$$L_T(w_T^*) - \min_{w \in \Delta_d} \sum_{t \in [T]} \ell_t(w) \leq \lambda(d-1) \log\left(2e \frac{T+\lambda d}{\lambda}\right). \tag{93}$$

Combining (90)–(93) yields the result. \square

B.2 Proof of Proposition 2.2

On the one hand, the mere feasibility of $\bar{\phi}_t$ in surrogate problem (27) implies, through (26), that

$$F_{t-1}[\bar{\phi}_t] \leq \bar{F}_{t-1}[\bar{\phi}_t]. \tag{94}$$

Now, let us define $(\underline{w}, \underline{\phi}_t)$ as the optimal solution to the following minimization problem:

$$\min_{\substack{\hat{w} \in \Delta_d, \mathbb{E}_\phi[w] = \hat{w}, \\ \phi \in \text{Supp}(\mathcal{E}_{t-1, 1/2}(\hat{w}))}} \underbrace{L_{t-1}(\hat{w}) + \frac{1}{5} \text{tr}(\nabla^2 L_{t-1}(\hat{w}) \text{Cov}[\phi]) - \mu \text{Ent}[\phi]}_{:= \underline{F}_{t-1}[\phi]}. \tag{95}$$

Moreover, let $\tilde{\phi}_t(\cdot) := 2\underline{\phi}_t(\underline{w} + 2(\cdot - \underline{w}))$, i.e. $\tilde{\phi}_t$ is the distribution of $\underline{w} + \frac{1}{2}(w - \underline{w})$ with $w \sim \underline{\phi}_t$. Then $\mathbb{E}_{w \sim \tilde{\phi}_t}[w] = \underline{w}$ and also $\tilde{\phi}_t \in \text{Supp}(\mathcal{E}_{t-1, 1/4}(\underline{w}))$ —in other words, $(\underline{w}, \tilde{\phi}_t)$ is feasible in (27), so

$$\bar{F}_{t-1}[\bar{\phi}_t] \leq \bar{F}_{t-1}[\tilde{\phi}_t]. \tag{96}$$

Now: observe that $\text{Cov}[\tilde{\phi}_t] = \frac{1}{4} \text{Cov}[\underline{\phi}_t]$ and $\text{Ent}[\tilde{\phi}_t] = \text{Ent}[\underline{\phi}_t] - (d-1) \log(2)$, whence (26) gives

$$\bar{F}_{t-1}[\tilde{\phi}_t] \leq \underline{F}_{t-1}[\underline{\phi}_t] + \mu(d-1) \log(2). \tag{97}$$

Finally, let (w^o, ϕ_t^o) be the optimal solution to the optimization problem

$$\min_{\substack{\hat{w} \in \Delta_d, \mathbb{E}_\phi[w] = \hat{w}, \\ \phi \in \text{Supp}(\mathcal{E}_{t-1, 1/2}(\hat{w}))}} F_{t-1}[\phi]. \tag{98}$$

Since (w^o, ϕ_t^o) is also a feasible solution to (95), we have that $\underline{F}_{t-1}[\underline{\phi}_t] \leq \underline{F}_{t-1}[\phi_t^o] \leq F_{t-1}[\phi_t^o]$. In combination with (94) and (96)–(98) this results in

$$F_{t-1}[\bar{\phi}_t] \leq \min_{\substack{\hat{w} \in \Delta_d, \mathbb{E}_{\phi}[w] = \hat{w}, \\ \phi \in \text{Supp}(\mathcal{E}_{t-1, 1/2}(\hat{w}))}} F_{t-1}[\phi] + \mu(d-1) \log(2). \quad (99)$$

Thus, to prove the proposition it suffices to bound the minimum in (98) from above in terms of the minimum in (23). To this end, recall that ϕ_t , as defined in (21), is the optimal solution to (23); let $w^* = \text{argmin}_{w \in \Delta_d} L_{t-1}(w)$, and let ϕ_t^{trc} be the truncation of ϕ_t to $\mathcal{E}_{t-1, 1/8}(w^*)$ —in other words,

$$\phi_t^{\text{trc}}(w) = \frac{\exp\left(-\frac{1}{\mu} L_{t-1}(w)\right)}{\int_{\mathcal{E}_{t-1, 1/8}(w^*)} \exp\left(-\frac{1}{\mu} L_{t-1}(w')\right) dw'}, \quad w \in \mathcal{E}_{t-1, 1/8}(w^*). \quad (100)$$

Clearly, ϕ_t^{trc} is supported on $\mathcal{E}_{t-1, 1/2}(\hat{w}^{\text{trc}})$ where $\hat{w}^{\text{trc}} = \mathbb{E}_{\phi_t^{\text{trc}}}[w]$ is its expectation—in other words, the pair $(\hat{w}^{\text{trc}}, \phi_t^{\text{trc}})$ is feasible in (98). Indeed, for any $w \in \mathcal{E}_{t-1, 1/8}(w^*)$ we have that

$$\begin{aligned} \|w - \hat{w}^{\text{trc}}\|_{\nabla^2 L_{t-1}(\hat{w}^{\text{trc}})} &\leq 2\|w - \hat{w}^{\text{trc}}\|_{\nabla^2 L_{t-1}(w^*)} \\ &\leq 2\|w - w^*\|_{\nabla^2 L_{t-1}(w^*)} + 2\|\hat{w}^{\text{trc}} - w^*\|_{\nabla^2 L_{t-1}(w^*)} \leq 1/2, \end{aligned}$$

where we used that function $L_{t-1}^\lambda(v)$ is 1-self-concordant whenever $\lambda \geq 1$ —see Corollary A.1 in Appendix A.3 for the proof of this result—and applied Lemma A.9. By (99), this implies that

$$F_{t-1}[\bar{\phi}_t] \leq F_{t-1}[\phi_t^{\text{trc}}] + \mu(d-1) \log(2). \quad (101)$$

As such, it remains to show that replacing ϕ_t with ϕ_t^{trc} does not lead to a dramatic increase of F_{t-1} .

Lemma B.1. *For $\phi_t, F_{t-1}, \phi_t^{\text{trc}}$ as in (21)–(23) and (100), the following holds as long as $\lambda \geq 1$:*

$$F_{t-1}[\phi_t^{\text{trc}}] \leq F_{t-1}[\phi_t] + 1.5\mu(d-1) \log(T + \lambda d) + 3.2\mu(d+1) + 0.1.$$

This lemma, proved in the next section, gives the desired result via (101), since $\log(2) < 0.7$. \square

B.3 Proof of Lemma B.1

For the sake of generality we shall consider truncation to $\mathcal{E}_{t-1, r}(w^*)$ with arbitrary $r \leq 1/2$, and put $r = 1/8$ post-hoc. By duality between the negative entropy and log-partition function (cf. (23)),

$$\begin{aligned} F_{t-1}[\phi_t] &= -\mu \log \int_{\Delta_d} \exp\left(-\frac{1}{\mu} L_{t-1}(w)\right) dw, \\ F_{t-1}[\phi_t^{\text{trc}}] &= -\mu \log \int_{\mathcal{E}_{t-1, r}(w^*)} \exp\left(-\frac{1}{\mu} L_{t-1}(w)\right) dw. \end{aligned}$$

Here the second integral is smaller: $\mathcal{E}_{t-1, r}(w^*) \subseteq \Delta_d$ by Lemma A.8 and Corollary A.1. Our goal is to show the reverse inequality up to the appropriate error term. To this end, we first note that

$$\exp\left(\frac{1}{\mu} (F_{t-1}[\phi_t^{\text{trc}}] - F_{t-1}[\phi_t])\right) = \frac{\int_{\Delta_d} \exp\left(-\frac{1}{\mu} [L_{t-1}(w) - L_{t-1}(w^*)]\right) dw}{\int_{\mathcal{E}_{t-1, r}(w^*)} \exp\left(-\frac{1}{\mu} [L_{t-1}(w) - L_{t-1}(w^*)]\right) dw}. \quad (102)$$

Since $L_{t-1}(w) \geq L_{t-1}(w^*)$, the numerator can be upper-bounded by the surface area of Δ_d :

$$\int_{\Delta_d} \exp\left(-\frac{1}{\mu}[L_{t-1}(w) - L_{t-1}(w^*)]\right) dw \leq \text{Area}(\Delta_d) = \sqrt{d}\text{Vol}(\mathbf{\Delta}_{d-1}) = \frac{\sqrt{d}}{\Gamma(d)}.$$

By the first-order optimality conditions $\langle \nabla L_{t-1}(w^*), w - w^* \rangle = 0$ for $w \in \Delta_d$, whence by Lemma A.8

$$L_{t-1}(w) - L_{t-1}(w^*) \leq \frac{4}{5} \|w - w^*\|_{\nabla^2 L_{t-1}(w^*)}^2 \quad \forall w \in \mathcal{E}_{t-1,r}(w^*),$$

where we used that $r \leq 1/2$. As such, we can lower-bound the denominator in (102) as follows:

$$\int_{\mathcal{E}_{t-1,r}(w^*)} \exp\left(-\frac{1}{\mu}[L_{t-1}(w) - L_{t-1}(w^*)]\right) dw \geq \exp\left(-\frac{4r^2}{5\mu}\right) r^{d-1} \text{Area}(\mathcal{E}_{t-1,1}(w^*)).$$

Now, as in Corollary A.1 let $L_{t-1}^{\mathbb{1}} : \mathbb{R}^{d-1} \rightarrow \overline{\mathbb{R}}$ be the reparametrized restriction of L_{t-1} to \mathbb{A}_d . Since $\det(A^\top A) = d$, we have that

$$\text{Area}(\mathcal{E}_{t-1,1}(w^*)) = \sqrt{d}\text{Vol}(\mathcal{E}_{t-1,1}^{\mathbb{1}}(v^*))$$

where v^* is such that $Av^* + e_d = w^*$, and ellipsoid $\mathcal{E}_{t-1,r}^{\mathbb{1}}(v^*)$ similarly corresponds to $\mathcal{E}_{t-1,r}(w^*)$:

$$\mathcal{E}_{t-1,r}^{\mathbb{1}}(v^*) := \{v \in \mathbb{R}^{d-1} : \|v - v^*\|_{\nabla^2 L_{t-1}^{\mathbb{1}}(v^*)} \leq r\} \subset \mathbb{R}^{d-1}.$$

Using the expression for the volume of an ellipsoid in \mathbb{R}^{d-1} , we have that

$$\text{Vol}(\mathcal{E}_{t-1,1}^{\mathbb{1}}(v^*)) = \frac{\pi^{\frac{d-1}{2}}}{\Gamma(\frac{d+1}{2}) \sqrt{\det(\nabla^2 L_{t-1}^{\mathbb{1}}(\bar{v}))}}.$$

Moreover, since $\nabla^2 L_{t-1}^{\mathbb{1}}(v^*) = A^\top \nabla^2 L_{t-1}(w^*) A$ and $A^\top A = \mathbf{I}_{d-1} + \mathbf{1}_{d-1} \mathbf{1}_{d-1}^\top$, so that $\|A\| = \sqrt{d}$,

$$\begin{aligned} \det(\nabla^2 L_{t-1}^{\mathbb{1}}(v^*)) &\leq \|\nabla^2 L_{t-1}^{\mathbb{1}}(v^*)\|^{d-1} \leq (d \|\nabla^2 L_{t-1}(w^*)\|)^{d-1} \leq \left(\frac{d(t-1+\lambda d)^3}{\lambda^2}\right)^{d-1} \\ &\leq (d(t-1+\lambda d)^3)^{d-1}; \end{aligned}$$

here we bounded $\|\nabla^2 L_{t-1}(w^*)\|$ by using that $\min_{i \in [d]} w^*[i] \geq \frac{\lambda}{t-1+\lambda d}$ (cf. Lemma A.1) and invoking estimate (43) used in step 1^o of the proof of Theorem 2.1. Returning to (102) we arrive at

$$\begin{aligned} F_{t-1}[\phi_t^{\text{trc}}] - F_{t-1}[\phi_t] &\leq \frac{3}{2} \mu (d-1) \log\left(\frac{T + \lambda d}{r^{2/3} \pi^{1/3}}\right) + \mu \log\left(\frac{\Gamma(\frac{d+1}{2}) d^{\frac{d-1}{2}}}{\Gamma(d)}\right) + \frac{4}{5} r^2 \\ &\leq \frac{3}{2} \mu (d-1) \log\left(\frac{T + \lambda d}{r}\right) + \mu (d+1) + \frac{4}{5} r^2. \end{aligned}$$

Here in the last step we did a series of simple estimates to bound the term depending solely on d . The result follows by plugging in $r = 1/8$. \square

B.4 Transition from (32) to (33)

Recall that $A^\top \nabla^2 L_{t-1}(\hat{w}) A = \nabla^2 L_{t-1}^1(\hat{v})$, cf. (11), and note that the Σ component of the objective gradient in (32) is $\frac{4}{5} \nabla^2 L_{t-1}^1(\hat{v}) - \frac{\mu}{2} \Sigma^{-1}$. The corresponding first-order optimality condition reads:

$$\text{tr} \left[(\Sigma - \hat{\Sigma}) \left(\nabla^2 L_{t-1}^1(\hat{v}) - \frac{5\mu}{8} \hat{\Sigma}^{-1} \right) \right] \geq 0 \quad \text{for all } 0 \preceq \Sigma \preceq \frac{1}{16} \nabla^2 L_{t-1}^1(\hat{v})^{-1}.$$

Plugging in $\hat{\Sigma} = \frac{1}{16} \nabla^2 L_{t-1}^1(\hat{v})^{-1}$ and dividing over $1 - 10\mu < 0$, the above condition translates into

$$\text{tr} \left[\Sigma \nabla^2 L_{t-1}^1(\hat{v}) - \frac{1}{16} \mathbf{I}_{d-1} \right] \leq 0 \quad \text{for all } 0 \preceq \Sigma \preceq \frac{1}{16} \nabla^2 L_{t-1}^1(\hat{v})^{-1}.$$

From the monotonicity of trace with respect to \preceq ordering, it follows that this condition is indeed satisfied. Plugging $\hat{\Sigma}$ into (32) and neglecting the constant terms, we finally arrive at (33). \square

B.5 Proof of Proposition 2.3

1^o. Recall that the optimization problem

$$\begin{aligned} \min_{\substack{\hat{w} \in \Delta_d, \quad 0 \preceq \Sigma \preceq \frac{1}{16} (A^\top \nabla^2 L_{t-1}(\hat{w}) A)^{-1}, \\ \phi^\mathbb{1} : \mathbb{A} \mathbb{E}_{v \sim \phi^\mathbb{1}}[v] + e_d = \hat{w}, \text{ Cov}[\phi^\mathbb{1}] = \Sigma}} \underbrace{L_{t-1}(\hat{w}) + \frac{4}{5} \text{tr} \left[\Sigma A^\top \nabla^2 L_{t-1}(\hat{w}) A \right] - \mu \text{Ent}[\phi^\mathbb{1}]}_{= \bar{F}_{t-1}[\phi] + \frac{\mu}{2} \log(d)} \end{aligned} \quad (103)$$

is equivalent to (32) in the natural sense: $(w_t, \Sigma_t) \in \Delta_d \times \mathbb{S}_+^{d-1}$ is optimal in (32) whenever $\mathcal{N}(v_t, \Sigma_t)$, with $v_t \in \Delta_{d-1}$ such that $A v_t + e_d = w_t$, is optimal in (103). Now, let $(\bar{w}, \bar{\phi}_t)$ and (w_t, Σ_t) be optimal in (27) and (32) correspondingly, so that $\mathcal{N}(v_t, \Sigma_t)$, with $v_t \in \Delta_{d-1}$ as above, is optimal in (103). Observe that $(\bar{w}, \bar{\phi}_t)$ is feasible in (103) due to (29). Since the objective in (103) amounts to $\bar{F}_{t-1}[\phi] + \frac{\mu}{2} \log(d)$ with $\phi \in \text{Supp}(\mathbb{A}_d)$ associated with $\phi^\mathbb{1}$ via (30), we have, for the distribution

$$g_t := \mathcal{N}(w_t, A \Sigma_t A^\top),$$

that

$$\bar{F}_{t-1}[g_t] \leq \bar{F}_{t-1}[\bar{\phi}_t]. \quad (104)$$

2^o. Next, let us bound $\bar{F}_{t-1}[g_t]$ from below. From our derivation in Appendix B.4 it follows that

$$\Sigma_t^{-1} = 16 A^\top \nabla^2 L_{t-1}(w_t) A, \quad (105)$$

therefore the level sets of g_t are precisely the Dikin ellipses $\mathcal{E}_{t-1,r}(w_t)$ with various $r \geq 0$, cf (24). Now, define g_t^{trc} as the truncation of g_t to the specific Dikin ellipse $\mathcal{E}_{t-1,1/2}(w_t)$ —in other words,

$$g_t^{\text{trc}}(w) = \frac{g_t(w)}{\int_{\mathcal{E}_{t-1,1/2}(w_t)} g_t(w') dw'}, \quad w \in \mathcal{E}_{t-1,1/2}(w_t). \quad (106)$$

Since $\mathbb{E}_{g_t^{\text{trc}}}[w] = w_t$ and $g_t^{\text{trc}} \in \text{Supp}(\mathcal{E}_{t-1,1/2}(w_t))$, the bounds in (26) are valid for $\phi = g_t^{\text{trc}}$, and so

$$F_{t-1}[g_t^{\text{trc}}] \leq \bar{F}_{t-1}[g_t^{\text{trc}}]. \quad (107)$$

3^o. Let us now bound $\bar{F}_{t-1}[g_t^{\text{trc}}]$ from above in terms of $\bar{F}_{t-1}[g_t]$. To this end, we first observe that

$$\text{Cov}[g_t^{\text{trc}}] \preceq \text{Cov}[g_t]. \quad (108)$$

Indeed: truncating g_t can also be seen as transporting the mass from each ellipse $\mathcal{E}_{t-1,s}(w_t)$, $s > 1/2$, to $\mathcal{E}_{t-1,r}(w_t)$ for some $r = r(s) \leq 1/2$. As such, it only remains to control the change of entropy because of truncation. The following lemma, proved in the next section, allows us to do so.

Lemma B.2. *Given $\hat{v} \in \mathbb{R}^n$, $\Sigma \in \mathbb{S}_{++}^n$, let $\mathbf{g}(v)$ be the density of $\mathcal{N}(\hat{v}, \Sigma)$. Consider the density*

$$\mathbf{g}^{(r)}(v) = \frac{\mathbf{g}(v)}{\int_{\|v'-\hat{v}\|_{\Sigma^{-1}} \leq r} \mathbf{g}(v') dv'}, \quad \forall v \in \mathbb{R}^n : \|v - \hat{v}\|_{\Sigma^{-1}} \leq r,$$

i.e. the truncation of \mathbf{g} to the covariance ellipsoid $\{v \in \mathbb{R}^n : \|v - \hat{v}\|_{\Sigma^{-1}} \leq r\}$ of radius $r > 0$. Then

$$\text{Ent}[\mathbf{g}] - \text{Ent}[\mathbf{g}^{(r)}] \leq \frac{n}{2} \log \left(\frac{2e(n+1)}{r^2} \right) + \frac{r^2}{2}.$$

While this bound becomes vacuous when $r \rightarrow \infty$, for our purposes it suffices, as in our case $r < 1$. Namely, we apply the lemma to $\mathcal{N}(v_t, \Sigma_t)$ in the role of \mathbf{g} , so that its truncation to $\mathcal{E}_{t-1,1/2}^{\mathbf{1}}(v_t)$, cf. (106), corresponds to $\mathbf{g}^{(1/8)}$, cf. (105). This results in

$$\text{Ent}[g_t] - \text{Ent}[g_t^{\text{trc}}] \leq \frac{d-1}{2} \log(128ed) + \frac{1}{128},$$

where we applied the identity $\text{Ent}[\phi^{\mathbf{1}}] = \text{Ent}[\phi] + \frac{1}{2} \log(d)$ to $\phi \in \{g_t, g_t^{\text{trc}}\}$. In view of (108) we get

$$\bar{F}_{t-1}[g_t^{\text{trc}}] \leq \bar{F}_{t-1}[g_t] + \frac{\mu(d-1)}{2} \log(128ed) + \frac{\mu}{128}.$$

The desired result follows by combining this with (107), (104), and Proposition 2.2. \square

B.6 Proof of Lemma B.2

Recall that $\text{Ent}[\mathbf{g}] = \frac{1}{2} \log \det(\Sigma) + \frac{n}{2} \log(2\pi) + \frac{n}{2}$. On the other hand, by rewriting $\mathbf{g}^{(r)}$ in the form

$$\mathbf{g}^{(r)}(v) = \frac{1}{\sqrt{(2\pi)^n C_r(\Sigma)}} \exp \left(-\frac{1}{2} \|v - \hat{v}\|_{\Sigma^{-1}}^2 \right)$$

where $C_r(\Sigma)$ does not depend on \hat{v} , we estimate $\text{Ent}[\mathbf{g}^{(r)}] \geq \frac{1}{2} \log[C_r(\Sigma)] + \frac{n}{2} \log(2\pi)$ and arrive at

$$\text{Ent}[\mathbf{g}^{(r)}] - \text{Ent}[\mathbf{g}] \geq \frac{1}{2} \log[C_r(\Sigma)] - \frac{1}{2} \log \det(\Sigma) - \frac{n}{2}. \quad (109)$$

Now, observe that, putting $\hat{v} = 0$ w.l.o.g.,

$$C_r(\Sigma)^{1/2} = \frac{1}{(2\pi)^{n/2}} \int_{\|v\|_{\Sigma^{-1}} \leq r} \exp \left(-\frac{1}{2} \|v\|_{\Sigma^{-1}}^2 \right) dv \geq \frac{e^{-r^2/2}}{(2\pi)^{n/2}} \det(\Sigma)^{1/2} \frac{\pi^{n/2} r^n}{\Gamma(\frac{n}{2} + 1)}$$

where we used the expression $\frac{\pi^{n/2}}{\Gamma(\frac{n}{2} + 1)}$ for the volume of the unit ℓ_2 -ball in \mathbb{R}^n . Whence we arrive at

$$\begin{aligned} \frac{1}{2} \log C_r(\Sigma) - \frac{1}{2} \log \det(\Sigma) &\geq -\frac{r^2}{2} - \frac{n}{2} \log(2) + n \log(r) - \log \left[\Gamma \left(\frac{n}{2} + 1 \right) \right] \\ &\geq -\frac{r^2}{2} - \frac{n}{2} \log(2) + n \log(r) - \frac{n}{2} \log(n+1), \end{aligned}$$

and the desired result follows via (109). \square

C Differential properties of leverage scores and barrier functions

In this section, we recall some differential properties of the Gram matrix $\mathbf{\Pi}(w)$, and barrier functions $V_t(w), P_t(w)$ used in VB-FTRL and in its analysis. These results are adaptations or easy corollaries of those in [45, 3]; however, the existing results in these works can only be applied in the special case $\lambda = \mu = 1$, where $V_t^\perp(\cdot)$ and $P_t^\perp(\cdot)$ become, respectively, the volumetric and “hybrid” barriers (in the terminology of [3]) for the set of linear constraints corresponding to the set of vectors $\{A^\top x_\tau, \tau \in [t]^+\}$ in the notation defined in (46)–(50).⁸ Formal reduction to this special case from $\lambda \geq 1, \mu > 0$ proved to be elusive. For this reason, and to make the paper self-contained, here we reprove these results assuming only that $\lambda, \mu > 0$, and in some cases that $\lambda \geq 1$.

Preliminaries and notation. To streamline the exposition, in this section we fix $t \in [T]$ and use the notation defined in (46)–(50), which we are now repeating here for convenience. For any $\tau \in [t]^+$:

$$\begin{aligned} \mathbf{H}_t(w) &:= A^\top \nabla^2 L_t(w) A, & \nabla_\tau(w) &:= A^\top \nabla \ell_\tau(w), \\ \hat{\mathbf{H}}_t &:= \mathbf{H}_t(w_t), & \hat{\nabla}_\tau &:= \nabla_\tau, \end{aligned} \quad (110)$$

where $[t] \cup \{-1, \dots, -d\}$ is the extended index set (with $\tau < 0$ corresponding to $x_\tau := e_{-\tau}$). Besides,

$$\begin{aligned} \lambda_\tau &:= \lambda^{\mathbf{1}\{\tau < 0\}}, \\ \text{so that } \mathbf{H}_t(w) &= \sum_{\tau \in [t]^+} \lambda_\tau \nabla_\tau(w) \nabla_\tau^\top(w) \quad \text{and} \quad \hat{\mathbf{H}}_t = \sum_{\tau \in [t]^+} \lambda_\tau \hat{\nabla}_\tau \hat{\nabla}_\tau^\top. \end{aligned} \quad (111)$$

We let $\mathbf{\Pi}(w) \in \mathbb{S}_+^{t+d}$ be the Gram matrix with rows and columns indexed over $[t]^+$, and with entries

$$\pi_{\tau, \nu}(w) := \sqrt{\lambda_\tau \lambda_\nu} \langle \nabla_\tau(w), \nabla_\nu(w) \rangle_{\mathbf{H}_t(w)^{-1}}. \quad (112)$$

We also define $\hat{\mathbf{\Pi}} := \mathbf{\Pi}(w_t)$, whose entries are thus

$$\hat{\pi}_{\tau, \nu} := \sqrt{\lambda_\tau \lambda_\nu} \langle \hat{\nabla}_\tau, \hat{\nabla}_\nu \rangle_{\hat{\mathbf{H}}_t^{-1}}. \quad (113)$$

For brevity, we refer to the diagonal entries of both these matrices with a single index. Note that

$$\pi_\tau(w) = \frac{\lambda_\tau}{(x_\tau^\top w)^2} \|A^\top x_\tau\|_{\mathbf{H}_t(w)^{-1}}^2. \quad (114)$$

C.1 Gram matrix $\mathbf{\Pi}(w)$ and leverage scores

Lemma C.1. *For any $t \in [T]$, $\tau, \nu \in [t]^+$, and $w \in \mathbb{R}_{++}^d$, one has the following estimate for $\pi_{\tau, \nu}$:*

$$|\pi_{\tau, \nu}(w)| \leq \sqrt{\frac{1}{(1 + \lambda)^{\mathbf{1}\{\tau > 0\} + \mathbf{1}\{\nu > 0\}}}}. \quad (115)$$

⁸The polyhedron defined by this set of constraints is Δ_{d-1} , same as without x_τ 's for $\tau \in [t]$. But functions L_t, V_t, P_t do depend on these x_τ 's, so their names adopted in [2]—e.g. “volumetric barrier” for V_t —should be used with caution.

Proof. By (112) and Cauchy-Schwarz, $|\pi_{\tau,\nu}(w)| \leq \sqrt{\pi_\tau(w)\pi_\nu(w)}$ for all $\tau, \nu \in [t]^+$. Clearly, for $i \in [d]$ one has $\nabla^2 L_t(w) \succcurlyeq \lambda \nabla^2 \ell_{-i}(w)$, whence $\pi_\tau(w) \leq 1$ for $\tau < 0$. Similarly, but also using Lemma A.3,

$$\nabla^2 L_t(w) \succcurlyeq \nabla^2 \ell_\tau(w) + \lambda \nabla^2 R(w) \succcurlyeq (1 + \lambda) \nabla^2 \ell_\tau(w), \quad \forall \tau \in [t].$$

As a result, $\pi_\tau \leq \frac{1}{1+\lambda}$ for $\tau \in [t]$. The result follows. \square

Lemma C.2 (Generalization of [45, Claim 3]). *For $t \in [T]$, $\tau_1, \tau_2 \in [t]^+$, and $w \in \mathbb{R}_{++}^d$, one has:*

$$\pi_{\tau_1, \tau_2}(w) = \sum_{\nu \in [t]^+} \pi_{\tau_1, \nu}(w) \pi_{\nu, \tau_2}(w). \quad (116)$$

In other words, $\mathbf{\Pi}(w)$ is a projection matrix. In particular, $\pi_\tau(w) = \sum_{\nu \in [t]^+} \pi_{\tau, \nu}(w)^2$ for $\tau \in [t]^+$.

Proof. Using the abridged notation defined in (110), and omitting the dependency on w for brevity,

$$\begin{aligned} \pi_{\tau_1, \tau_2} &\stackrel{(112)}{=} \sqrt{\lambda_{\tau_1} \lambda_{\tau_2}} \nabla_{\tau_1}^\top \mathbf{H}_t^{-1} \nabla_{\tau_2} = \sqrt{\lambda_{\tau_1} \lambda_{\tau_2}} \nabla_{\tau_1}^\top \mathbf{H}_t^{-1} \mathbf{H}_t \mathbf{H}_t^{-1} \nabla_{\tau_2} \\ &\stackrel{(111)}{=} \sqrt{\lambda_{\tau_1} \lambda_{\tau_2}} \nabla_{\tau_1}^\top \mathbf{H}_t^{-1} \left(\sum_{\nu \in [t]^+} \lambda_\nu \nabla_\nu \nabla_\nu^\top \right) \mathbf{H}_t^{-1} \nabla_{\tau_2} \\ &\stackrel{(112)}{=} \sum_{\nu \in [t]^+} \pi_{\tau_1, \nu} \pi_{\nu, \tau_2}. \quad \square \end{aligned}$$

Lemma C.3 (Gradient of a leverage score). *For $t \in [T]$, $\tau \in [t]^+$, and $w \in \mathbb{R}_{++}^d$, it holds that*

$$\frac{1}{2} \nabla \pi_\tau(w) = \pi_\tau(w) \nabla \ell_\tau(w) - \sum_{\nu \in [t]^+} \pi_{\tau, \nu}(w)^2 \nabla \ell_\nu(w). \quad (117)$$

Proof. This is a byproduct of an explicit formula for $\nabla^2 V_t(w)$ in the proof of Lemma C.4. \square

C.2 Gradient and Hessian of the volumetric barrier

Following [45] and [3], we shall now derive an explicit formula for the gradient of function $V_t(\cdot)$ on Δ_d , cf. (16), estimate its Hessian, and verify the expression for $\nabla \pi_\tau(w)$ in Lemma C.3, cf. (123)–(124). The lower bound on $\nabla^2 V_t(w)$ in (119) shows, in particular, that $V_t(\cdot)$ is strictly convex.

Lemma C.4 ([45, Lemmas 1–3]). *For $t \in [T]$ and $w \in \mathbb{R}_{++}^d$, the gradient of $V_t(w)$, cf. (16), is*

$$\nabla V_t(w) = \sum_{\tau \in [t]^+} \pi_\tau(w) \nabla \ell_\tau(w), \quad (118)$$

and its Hessian satisfies

$$\mathbf{Q}_t(w) \preceq \nabla^2 V_t(w) \preceq 3\mathbf{Q}_t(w) \quad \text{where} \quad \mathbf{Q}_t(w) = \sum_{\tau \in [t]^+} \pi_\tau(w) \nabla^2 \ell_\tau(w). \quad (119)$$

Remark C.1. A weaker form of the upper bound in (119), namely $\nabla^2 V_t(w) \preceq 5\mathbf{Q}_t(w)$, was first obtained by Vaidya in [45]; later on, [2] improved it to $\nabla^2 V_t(w) \preceq 3\mathbf{Q}_t(w)$. (In both cases, it was also assumed that $\lambda = 1$.) For the sake of clarity, here we prove the bound with factor 5. The improved bound follows by applying Schur's product theorem (see e.g. [4]) to the Hadamard-product representation of the matrix $\mathbf{S}_t(w) = \frac{3}{2}\mathbf{Q}_t(w) - \frac{1}{2}\nabla^2 V_t(w)$, see (131) in the proof of Lemma C.5.

Proof. **1°.** Note that the entries of $\mathbf{H}_t(w) = A^\top \nabla^2 L_t(w) A$, for $i, j \in [d-1]$, are given by

$$[\mathbf{H}_t(w)]_{ij} = \sum_{\tau \in [t]^+} \frac{\lambda_\tau}{(x_\tau^\top w)^2} [A^\top x_\tau]_i [A^\top x_\tau]_j.$$

As such,

$$\nabla([\mathbf{H}_t(w)]_{ij}) = - \sum_{\tau \in [t]^+} \frac{2\lambda_\tau}{(x_\tau^\top w)^3} [A^\top x_\tau]_i [A^\top x_\tau]_j x_\tau. \quad (120)$$

Using this formula, the composition rule now allows us to verify (118) by computing $\nabla V_t(w)$ directly:

$$\begin{aligned} \nabla V_t(w) &= \frac{1}{2} \sum_{i,j \in [d-1]} [\mathbf{H}_t(w)^{-1}]_{ij} \nabla([\mathbf{H}_t(w)]_{ij}) \\ &\stackrel{(120)}{=} - \sum_{\tau \in [t]^+} \frac{\lambda_\tau}{(x_\tau^\top w)^3} \left(\sum_{i,j \in [d-1]} [A^\top x_\tau]_i [\mathbf{H}_t(w)^{-1}]_{ij} [A^\top x_\tau]_j \right) x_\tau \\ &= - \sum_{\tau \in [t]^+} \frac{\lambda_\tau \|A^\top x_\tau\|_{\mathbf{H}_t(w)^{-1}}^2}{(x_\tau^\top w)^3} x_\tau = \sum_{\tau \in [t]^+} \pi_\tau(w) \nabla \ell_\tau(w). \end{aligned}$$

2°. Our next goal is to derive the following explicit formula:

$$\nabla^2 V_t(w) = 3\mathbf{Q}_t(w) - 2\mathbf{S}_t(w) \quad \text{where} \quad \mathbf{S}_t(w) := \sum_{\tau, \nu \in [t]^+} \pi_{\tau, \nu}(w)^2 \nabla \ell_\tau(w) \nabla \ell_\nu(w)^\top. \quad (121)$$

But before, let us show how (119) follows from it. Indeed, recall that $\pi_\tau(w) = \sum_{\nu \in [t]^+} \pi_{\tau, \nu}(w)^2$, cf. Lemma C.2. Whence from (121) and the definition of $\mathbf{Q}_t(w)$, cf. (119), we obtain

$$\begin{aligned} \nabla^2 V_t(w) - \mathbf{Q}_t(w) &= 2\mathbf{Q}_t(w) - 2\mathbf{S}_t(w) \\ &= \sum_{\tau, \nu \in [t]^+} \pi_{\tau, \nu}(w)^2 \left(2\nabla^2 \ell_\tau(w) - 2\nabla \ell_\tau(w) \nabla \ell_\nu(w)^\top \right) \\ &= \sum_{\tau, \nu \in [t]^+} \pi_{\tau, \nu}(w)^2 \left(\nabla \ell_\tau(w) - \nabla \ell_\nu(w) \right) \left(\nabla \ell_\tau(w) - \nabla \ell_\nu(w) \right)^\top \succeq 0, \end{aligned}$$

the lower bound in (119). We obtain the upper bound $\nabla^2 V_t(w) \preceq 5\mathbf{Q}_t(w)$ in a similar fashion:

$$\begin{aligned} 5\mathbf{Q}_t(w) - \nabla^2 V_t(w) &= 2\mathbf{Q}_t(w) + 2\mathbf{S}_t(w) \\ &= \sum_{\tau, \nu \in [t]^+} \pi_{\tau, \nu}(w)^2 \left(\nabla \ell_\tau(w) + \nabla \ell_\nu(w) \right) \left(\nabla \ell_\tau(w) + \nabla \ell_\nu(w) \right)^\top \succeq 0. \end{aligned}$$

3^o. It remains to verify (121). First, direct differentiation of (118) shows that

$$\nabla^2 V_t(w) = \mathbf{Q}_t(w) + \mathbf{R}_t(w) \quad \text{where} \quad \mathbf{R}_t(w) := \sum_{\tau \in [t]^+} \nabla \pi_\tau(w) \nabla \ell_\tau(w)^\top, \quad (122)$$

granted that $\mathbf{R}_t(w)$ is symmetric. Thus, to get (121) we must show that $\mathbf{R}_t(w) = 2\mathbf{Q}_t(w) - 2\mathbf{S}_t(w)$. To this end, observe that, by (110), (114), and the product rule,

$$\begin{aligned} \nabla \pi_\tau(w) &= \frac{\lambda_\tau}{(x_\tau^\top w)^2} \nabla \left(\|A^\top x_\tau\|_{\mathbf{H}_t(w)^{-1}}^2 \right) - \frac{2\lambda_\tau \|A^\top x_\tau\|_{\mathbf{H}_t(w)^{-1}}^2}{(x_\tau^\top w)^3} x_\tau \\ &= \frac{\lambda_\tau}{(x_\tau^\top w)^2} \nabla \left(\|A^\top x_\tau\|_{\mathbf{H}_t(w)^{-1}}^2 \right) + 2\pi_\tau(w) \nabla \ell_\tau(w). \end{aligned} \quad (123)$$

Furthermore, by Lemma A.5 and the composition rule,

$$\begin{aligned} \nabla \left(\|A^\top x_\tau\|_{\mathbf{H}_t(w)^{-1}}^2 \right) &= - \sum_{i,j \in [d-1]} \left[\mathbf{H}_t(w)^{-1} A^\top x_\tau x_\tau^\top A \mathbf{H}_t(w)^{-1} \right]_{ij} \nabla ([\mathbf{H}_t(w)]_{ij}) \\ &\stackrel{(120)}{=} \sum_{\nu \in [t]^+} \frac{2\lambda_\nu}{(x_\nu^\top w)^3} \left(\sum_{i,j \in [d-1]} [A^\top x_\nu]_i \left[\mathbf{H}_t(w)^{-1} A^\top x_\tau x_\tau^\top A \mathbf{H}_t(w)^{-1} \right]_{ij} [A^\top x_\nu]_j \right) x_\nu \\ &= \sum_{\nu \in [t]^+} \frac{2\lambda_\nu}{(x_\nu^\top w)^3} \left\langle A^\top x_\nu, A^\top x_\tau \right\rangle_{\mathbf{H}_t(w)^{-1}}^2 x_\nu \\ &= - \sum_{\nu \in [t]^+} \frac{2\lambda_\nu}{(x_\nu^\top w)^2} \left\langle A^\top x_\nu, A^\top x_\tau \right\rangle_{\mathbf{H}_t(w)^{-1}}^2 \nabla \ell_\nu(w), \end{aligned}$$

whence

$$\begin{aligned} \frac{\lambda_\tau}{(x_\tau^\top w)^2} \nabla \left(\|A^\top x_\tau\|_{\mathbf{H}_t(w)^{-1}}^2 \right) &= -2 \sum_{\nu \in [t]^+} \lambda_\tau \lambda_\nu \left\langle A^\top \nabla \ell_\tau(w), A^\top \nabla \ell_\nu(w) \right\rangle_{\mathbf{H}_t(w)^{-1}}^2 \nabla \ell_\nu(w) \\ &= -2 \sum_{\nu \in [t]^+} \pi_{\tau,\nu}(w)^2 \nabla \ell_\nu(w). \end{aligned} \quad (124)$$

Combining (122)–(124), we verify the identity $\mathbf{R}_t(w) = 2\mathbf{Q}_t(w) - 2\mathbf{S}_t(w)$, and hence also (121). \square

C.3 Self-concordance of the hybrid barrier

The following result can be understood as an extension of Corollary A.1 from $L_t(\cdot)$ to P_t with $\mu > 0$.

Proposition C.1. *Assume $\lambda \geq 1$ and $\mu > 0$, then the restriction $P_t^\mathbb{1} : \mathbb{R}^{d-1} \rightarrow \overline{\mathbb{R}}$ of P_t to \mathbb{A}_d , i.e.*

$$P_t^\mathbb{1}(v) = P_t(Av + e_d),$$

cf. (17) and (8)–(10), is 21-self-concordant according to Definition 1 with domain $\text{int}(\Delta_{d-1})$, cf. (7).

Proof. The requirements on the domain are verified in the same way as in the proof of Corollary A.1. Thus, it remains to show that when λ, μ are as in the premise, for any $u \in \mathbb{R}^d$ and $w \in \Delta_d$ one has

$$|\nabla^3 P_t(w)[u, u, u]| \leq 42 \left(\nabla^2 P_t(w)[u, u] \right)^{3/2}. \quad (125)$$

Indeed, this implies that $|\nabla^3 P_t^\mathbb{1}(v)[u^\mathbb{1}, u^\mathbb{1}, u^\mathbb{1}]| \leq 42(\nabla^2 P_t^\mathbb{1}(v)[u^\mathbb{1}, u^\mathbb{1}])^{3/2}$ for $v \in \Delta_{d-1}$, $u^\mathbb{1} \in \mathbb{R}^{d-1}$ by affine invariance, i.e. since for $w = Av + e_d$ and $u = Au^\mathbb{1}$ one has $\nabla^2 P_t^\mathbb{1}(v)[u^\mathbb{1}, u^\mathbb{1}] = \nabla^2 P_t(w)[u, u]$ and $\nabla^3 P_t^\mathbb{1}(v)[u^\mathbb{1}, u^\mathbb{1}, u^\mathbb{1}] = \nabla^3 P_t(w)[u, u, u]$. In turn, (125) can be derived from the following result.

Lemma C.5 (Adaptation of [3, Thm. 4.1]). *Let $w, \bar{w} \in \mathbb{R}_{++}^d$ be such that $\delta_t(w, \bar{w}) \leq 0.1$, where*

$$\delta_t(w, \bar{w}) = \max_{\tau \in [t]^+} \left| \frac{x_\tau^\top (\bar{w} - w)}{x_\tau^\top w} \right|.$$

Then for $\varepsilon_t(w, \bar{w}) = \frac{42\delta_t(w, \bar{w})}{[1 - \delta_t(w, \bar{w})]^4}$ one has that

$$\begin{aligned} -\varepsilon_t(w, \bar{w}) \nabla^2 V_t(w) &\preceq \nabla^2 V_t(\bar{w}) - \nabla^2 V_t(w) \preceq \varepsilon_t(w, \bar{w}) \nabla^2 V_t(w), \\ -\varepsilon_t(w, \bar{w}) \nabla^2 P_t(w) &\preceq \nabla^2 P_t(\bar{w}) - \nabla^2 P_t(w) \preceq \varepsilon_t(w, \bar{w}) \nabla^2 P_t(w). \end{aligned}$$

Now, fix $u \in \mathbb{R}^d$ and note that $\nabla^3 P_t(w)[u, u, u] = \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} (\nabla^2 P_t(w + \alpha u)[u, u] - \nabla^2 P_t(w)[u, u])$. On the other hand, when $|\alpha|$ is small enough we can apply Lemma C.5 to $\bar{w} = w + \alpha u$. Indeed:

$$\delta_t(w, w + \alpha u) = |\alpha| M_t(u, w) \quad \text{where} \quad M_t(u, w) := \max_{\tau \in [t]^+} \left| \frac{x_\tau^\top u}{x_\tau^\top w} \right|,$$

therefore $\delta_t(w, w + \alpha u) \rightarrow 0$ when $|\alpha| \rightarrow 0$ with fixed u, w . By Lemma C.5 we then conclude that

$$|\nabla^3 P_t(w)[u, u, u]| \leq \lim_{\alpha \rightarrow 0} \frac{42 \delta_t(w, w + \alpha u)}{|\alpha| [1 - \delta_t(w, w + \alpha u)]^4} \nabla^2 P_t(w)[u, u] = 42 M_t(u, w) \nabla^2 P_t(w)[u, u].$$

Finally, note that $\nabla^2 P_t(w)[u, u] \geq \nabla^2 L_t(w)[u, u] \geq M_t(u, w)^2 \min\{\lambda, 1\} = M_t(u, w)^2$ as $\lambda \geq 1$. \square

Proof of Lemma C.5. Our proof closely follows the one [3] while allowing for $\lambda \geq 1$ and $\mu > 0$. In the remainder of this section, we let $\delta := \delta_t(w, \bar{w})$ and also introduce some abbreviated notation:

$$\begin{aligned} \eta_\tau &:= x_\tau^\top w, & \bar{\eta}_\tau &:= x_\tau^\top \bar{w}, & \pi_\tau &:= \pi_\tau(w), & \bar{\pi}_\tau &:= \pi_\tau(\bar{w}), & \nabla_\tau &:= \nabla_\tau(w), & \bar{\nabla}_\tau &:= \nabla_\tau(\bar{w}), \\ \mathbf{H} &:= \mathbf{H}_t(w), & \bar{\mathbf{H}} &:= \mathbf{H}_t(\bar{w}), & \mathbf{Q} &:= \mathbf{Q}_t(w), & \bar{\mathbf{Q}} &:= \mathbf{Q}_t(\bar{w}), & \mathbf{S} &:= \mathbf{S}_t(w), & \bar{\mathbf{S}} &:= \mathbf{S}_t(\bar{w}). \end{aligned}$$

Lemma C.6. *Assuming that $\delta < 1$, for any $\tau \in [t]^+$ the following bounds hold:*

$$1 - \delta \leq \frac{\bar{\eta}_\tau}{\eta_\tau} \leq 1 + \delta, \quad \frac{1}{(1 + \delta)^2} \mathbf{H} \preceq \bar{\mathbf{H}} \preceq \frac{1}{(1 - \delta)^2} \mathbf{H}, \quad (126)$$

$$\left(\frac{1 - \delta}{1 + \delta} \right)^2 \leq \frac{\bar{\pi}_\tau}{\pi_\tau} \leq \left(\frac{1 + \delta}{1 - \delta} \right)^2, \quad \frac{(1 - \delta)^2}{(1 + \delta)^4} \mathbf{Q} \preceq \bar{\mathbf{Q}} \preceq \frac{(1 + \delta)^2}{(1 - \delta)^4} \mathbf{Q}. \quad (127)$$

Proof. The bounds for $\bar{\eta}$ are obvious from the definition of δ , and immediately imply those for $\bar{\mathbf{H}}$:

$$\bar{\mathbf{H}} = \sum_{\tau \in [t]^+} \frac{\lambda_\tau}{\bar{\eta}_\tau^2} A^\top x_\tau x_\tau^\top A \preceq \frac{1}{(1 - \delta)^2} \sum_{\tau \in [t]^+} \frac{\lambda_\tau}{\eta_\tau^2} A^\top x_\tau x_\tau^\top A = \frac{1}{(1 - \delta)^2} \mathbf{H},$$

and similarly for the corresponding lower bound. Furthermore, (126) implies

$$\frac{\bar{\pi}_\tau}{\pi_\tau} = \frac{\lambda_\tau \|\bar{\nabla}_\tau\|_{\bar{\mathbf{H}}^{-1}}^2}{\lambda_\tau \|\nabla_\tau\|_{\mathbf{H}^{-1}}^2} = \frac{\|\bar{\nabla}_\tau\|_{\bar{\mathbf{H}}^{-1}}^2}{\|\nabla_\tau\|_{\mathbf{H}^{-1}}^2} \leq (1 + \delta)^2 \frac{\|\bar{\nabla}_\tau\|_{\bar{\mathbf{H}}^{-1}}^2}{\|\nabla_\tau\|_{\mathbf{H}^{-1}}^2} = (1 + \delta)^2 \frac{\eta_\tau^2}{\bar{\eta}_\tau^2} \leq \left(\frac{1 + \delta}{1 - \delta} \right)^2,$$

and similarly for the corresponding lower bound. Finally, from the above results we conclude that

$$\bar{\mathbf{Q}} = \sum_{\tau \in [t]^+} \frac{\bar{\pi}_\tau}{\bar{\eta}_\tau^2} x_\tau x_\tau^\top \preceq \frac{(1+\delta)^2}{(1-\delta)^4} \sum_{\tau \in [t]^+} \frac{\pi_\tau}{\eta_\tau^2} x_\tau x_\tau^\top = \frac{(1+\delta)^2}{(1-\delta)^4} \mathbf{Q},$$

and similarly for the corresponding lower bound. \square

1^o. Returning to the proof of Lemma C.5, recall that

$$\nabla^2 P_t(\bar{w}) - \nabla^2 P_t(w) = \nabla^2 L_t(\bar{w}) - \nabla^2 L_t(w) + \mu [\nabla^2 V_t(\bar{w}) - \nabla^2 V_t(w)],$$

and let us consider the two differences in the right-hand side separately. For the first difference we observe that $\nabla^2 L_t(w)$ and $\nabla^2 L_t(\bar{w})$ satisfy the same relative bounds as \mathbf{H} and $\bar{\mathbf{H}}$ in (126), so

$$\nabla^2 L_t(\bar{w}) - \nabla^2 L_t(w) \preceq \left(\frac{1}{(1-\delta)^2} - 1 \right) \nabla^2 L_t(w) \preceq \frac{2\delta}{(1-\delta)^4} \nabla^2 L_t(w)$$

where we used that $0 \leq \delta \leq 1$. Similarly,

$$\nabla^2 L_t(w) - \nabla^2 L_t(\bar{w}) \preceq \left(1 - \frac{1}{(1+\delta)^2} \right) \nabla^2 L_t(w) \preceq \frac{3\delta}{(1-\delta)^4} \nabla^2 L_t(w).$$

Thus, to prove the lemma it suffices to show the first of the claimed inequalities, that is

$$-\frac{42\delta}{(1-\delta)^4} \nabla^2 V_t(w) \preceq \nabla^2 V_t(\bar{w}) - \nabla^2 V_t(w) \preceq \frac{42\delta}{(1-\delta)^4} \nabla^2 V_t(w). \quad (128)$$

2^o. To this end, we first observe that

$$\nabla^2 V_t(\bar{w}) - \nabla^2 V_t(w) \stackrel{(121)}{=} 3(\bar{\mathbf{Q}} - \mathbf{Q}) + 2(\bar{\mathbf{S}} - \mathbf{S}),$$

and Lemmas C.6–C.4 already allow us to control the difference $\bar{\mathbf{Q}} - \mathbf{Q}$, namely:

$$\begin{aligned} \bar{\mathbf{Q}} - \mathbf{Q} &\stackrel{(127)}{\preceq} \left(\frac{(1+\delta)^2}{(1-\delta)^4} - 1 \right) \mathbf{Q} = \frac{(3\delta - \delta^2)(2 - \delta + \delta^2)}{(1-\delta)^4} \mathbf{Q} \stackrel{(a)}{\preceq} \frac{6\delta}{(1-\delta)^4} \mathbf{Q}, \\ \mathbf{Q} - \bar{\mathbf{Q}} &\stackrel{(127)}{\preceq} \left(1 - \frac{(1-\delta)^2}{(1+\delta)^4} \right) \mathbf{Q} \stackrel{(b)}{\preceq} \left(\frac{(1+\delta)^2}{(1-\delta)^4} - 1 \right) \mathbf{Q} \preceq \frac{6\delta}{(1-\delta)^4} \mathbf{Q}. \end{aligned} \quad (129)$$

Here in (a) we used that $0 \leq \delta < 1$, and (b) holds since the function $\frac{(1+\delta)^2}{(1-\delta)^4} + \frac{(1-\delta)^2}{(1+\delta)^4}$ on $[0, 1)$ is minimized at $\delta = 0$. Let us now study the difference $\bar{\mathbf{S}} - \mathbf{S}$.

3^o. For brevity, we define $\mathbf{\Pi} := \mathbf{\Pi}_t(w)$ and $\bar{\mathbf{\Pi}} := \mathbf{\Pi}_t(\bar{w})$. We also define some auxiliary matrices:

- matrix $X \in \mathbb{R}^{d \times (t+d)}$ with x_τ in its τ^{th} column;
- diagonal matrices $\mathbf{\Lambda}, \mathbf{D}, \bar{\mathbf{D}}, \mathbf{Z} \in \mathbb{S}_{++}^{t+d}$ with $\lambda_\tau, \eta_\tau, \bar{\eta}_\tau, \frac{\eta_\tau}{\bar{\eta}_\tau}$ at τ^{th} diagonal position, respectively.

As earlier for $\mathbf{\Pi}_t(w)$, we use the convention that indexing is over $[t]^+$ in each dimension where the size of a matrix is $t+d$. Finally, let $\mathbf{A} \circ \mathbf{B}$ be the Hadamard (i.e. entrywise) product of two matrices with the same dimensions, and let $\mathbf{M}^{(2)} = \mathbf{M} \circ \mathbf{M}$. In terms of this matrix notation, we have that

$$\mathbf{Q} = \mathbf{X} \mathbf{D}^{-1} \mathbf{\Psi} \mathbf{D}^{-1} \mathbf{X}^\top, \quad (130)$$

$$\mathbf{S} = \mathbf{X} \mathbf{D}^{-1} \mathbf{\Pi}^{(2)} \mathbf{D}^{-1} \mathbf{X}^\top, \quad (131)$$

where $\mathbf{\Psi} = \text{Diag}(\mathbf{\Pi})$ is the diagonal matrix with π_τ at τ^{th} diagonal position. On the other hand,

$$\begin{aligned} \bar{\mathbf{S}} &= \mathbf{X} \bar{\mathbf{D}}^{-1} \bar{\mathbf{\Pi}}^{(2)} \bar{\mathbf{D}}^{-1} \mathbf{X}^\top \\ &= \mathbf{X} \mathbf{D}^{-1} \mathbf{Z} \bar{\mathbf{\Pi}}^{(2)} \mathbf{Z} \mathbf{D}^{-1} \mathbf{X}^\top. \end{aligned} \quad (132)$$

where in the last step we used that $\mathbf{Z} = \bar{\mathbf{D}}^{-1} \mathbf{D}$. In a similar way, $\mathbf{\Pi} = \mathbf{D}^{-1} \mathbf{\Lambda}^{1/2} \mathbf{X}^\top \mathbf{H}^{-1} \mathbf{X} \mathbf{\Lambda}^{1/2} \mathbf{D}^{-1}$ and $\bar{\mathbf{\Pi}} = \bar{\mathbf{D}}^{-1} \mathbf{\Lambda}^{1/2} \mathbf{X}^\top \bar{\mathbf{H}}^{-1} \mathbf{X} \mathbf{\Lambda}^{1/2} \bar{\mathbf{D}}^{-1} = \mathbf{Z} \mathbf{D}^{-1} \mathbf{\Lambda}^{1/2} \mathbf{X}^\top \bar{\mathbf{H}}^{-1} \mathbf{X} \mathbf{\Lambda}^{1/2} \mathbf{D}^{-1} \mathbf{Z}$, whence by (126) we get

$$(1 - \delta)^2 \mathbf{Z} \mathbf{\Pi} \mathbf{Z} \preceq \bar{\mathbf{\Pi}} \preceq (1 + \delta)^2 \mathbf{Z} \mathbf{\Pi} \mathbf{Z}. \quad (133)$$

Since \mathbf{Z} is diagonal, $(\mathbf{Z} \mathbf{\Pi} \mathbf{Z})^{(2)} = \mathbf{Z}^2 \mathbf{\Pi}^{(2)} \mathbf{Z}^2$. Moreover, recall that for two symmetric matrices \mathbf{A}, \mathbf{B} such that $\mathbf{A} \succeq \mathbf{B} \succeq 0$, it holds that $\mathbf{A}^{(2)} \succeq \mathbf{B}^{(2)}$.⁹ In combination with (133), these two facts imply

$$(1 - \delta)^4 \mathbf{Z}^2 \mathbf{\Pi}^{(2)} \mathbf{Z}^2 \preceq \bar{\mathbf{\Pi}}^{(2)} \preceq (1 + \delta)^4 \mathbf{Z}^2 \mathbf{\Pi}^{(2)} \mathbf{Z}^2.$$

Whence by (132) we conclude that

$$(1 - \delta)^4 \mathbf{X} \mathbf{D}^{-1} \mathbf{Z}^3 \mathbf{\Pi}^{(2)} \mathbf{Z}^3 \mathbf{D}^{-1} \mathbf{X}^\top \preceq \bar{\mathbf{S}} \preceq (1 + \delta)^4 \mathbf{X} \mathbf{D}^{-1} \mathbf{Z}^3 \mathbf{\Pi}^{(2)} \mathbf{Z}^3 \mathbf{D}^{-1} \mathbf{X}^\top. \quad (134)$$

We are now going to use identities (130)–(131) and inequality (134) to bound $\bar{\mathbf{S}} - \mathbf{S}$ in terms of \mathbf{Q} .

4^o. Defining $\mathbf{J} := \mathbf{\Psi}^{-1/2} \mathbf{\Pi}^{(2)} \mathbf{\Psi}^{-1/2}$ and $\mathbf{U} := \mathbf{X} \mathbf{D}^{-1} \mathbf{\Psi}^{1/2}$ we rewrite (130), (131), (134):

$$\begin{aligned} \mathbf{Q} &= \mathbf{U} \mathbf{U}^\top, \quad \mathbf{S} = \mathbf{U} \mathbf{J} \mathbf{U}^\top, \\ (1 - \delta)^4 \mathbf{U} \mathbf{Z}^3 \mathbf{J} \mathbf{Z}^3 \mathbf{U}^\top &\preceq \bar{\mathbf{S}} \preceq (1 + \delta)^4 \mathbf{U} \mathbf{Z}^3 \mathbf{J} \mathbf{Z}^3 \mathbf{U}^\top. \end{aligned} \quad (135)$$

We are now going to bound $\mathbf{M}_1 := (1 + \delta)^4 \mathbf{Z}^3 \mathbf{J} \mathbf{Z}^3 - \mathbf{J}$ and $\mathbf{M}_2 := \mathbf{J} - (1 - \delta)^4 \mathbf{Z}^3 \mathbf{J} \mathbf{Z}^3$ from above. Since $\mathbf{\Psi}, \mathbf{Z}$ commute as diagonal matrices, and by similarity, \mathbf{M}_1 has the same eigenvalues as

$$\mathbf{\Psi}^{-1/2} \mathbf{M}_1 \mathbf{\Psi}^{1/2} = \mathbf{\Psi}^{-1} \left((1 + \delta)^4 \mathbf{Z}^3 \mathbf{\Pi}^{(2)} \mathbf{Z}^3 - \mathbf{\Pi}^{(2)} \right).$$

The entries of the latter matrix are given by

$$[\mathbf{\Psi}^{-1/2} \mathbf{M}_1 \mathbf{\Psi}^{1/2}]_{\tau, \nu} = \frac{\pi_{\tau, \nu}^2}{\pi_\tau} \left((1 + \delta)^4 z_\tau^3 z_\nu^3 - 1 \right) \quad (136)$$

where $z_\tau = \frac{\eta_\tau}{\bar{\eta}_\tau}$ is the τ^{th} diagonal entry of \mathbf{Z} . By the Gershgorin circle theorem [18], each eigenvalue of the matrix $\mathbf{\Psi}^{-1/2} \mathbf{M}_1 \mathbf{\Psi}^{1/2}$ —and hence also of \mathbf{M}_1 —belongs to at least one of the segments

$$\left[[\mathbf{\Psi}^{-1/2} \mathbf{M}_1 \mathbf{\Psi}^{1/2}]_{\tau, \tau} - r_\tau, [\mathbf{\Psi}^{-1/2} \mathbf{M}_1 \mathbf{\Psi}^{1/2}]_{\tau, \tau} + r_\tau \right] \quad \text{with } r_\tau := \sum_{\nu \neq \tau} |[\mathbf{\Psi}^{-1/2} \mathbf{M}_1 \mathbf{\Psi}^{1/2}]_{\tau, \nu}|, \quad \tau \in [t]^+.$$

⁹This can be seen by writing $\mathbf{A}^{(2)} - \mathbf{B}^{(2)} = \mathbf{A} \circ (\mathbf{A} - \mathbf{B}) + (\mathbf{A} - \mathbf{B}) \circ \mathbf{A}$ and applying the Schur product theorem [4].

Whence by (136), (126):

$$\mathbf{M}_1 \preceq \max_{\tau \in [t]^+} \left\{ \frac{1}{\pi_\tau} \sum_{\nu \in [t]^+} \pi_{\tau, \nu}^2 \right\} \left(\frac{(1+\delta)^4}{(1-\delta)^6} - 1 \right) \mathbf{I} \stackrel{(116)}{\preceq} \left(\frac{(1+\delta)^4}{(1-\delta)^6} - 1 \right) \mathbf{I}.$$

Proceeding with \mathbf{M}_2 in a similar fashion as we did with \mathbf{M}_1 , we get

$$\mathbf{M}_2 \preceq \left(1 - \frac{(1-\delta)^4}{(1+\delta)^6} \right) \mathbf{I} \preceq \left(\frac{(1+\delta)^4}{(1-\delta)^6} - 1 \right) \mathbf{I}.$$

(The second step uses that the function $\frac{(1+\delta)^4}{(1-\delta)^6} + \frac{(1-\delta)^4}{(1+\delta)^6}$ increases on $[0, 1)$.) Finally, observe that

$$\frac{(1+\delta)^4}{(1-\delta)^6} - 1 = \frac{4\delta(1+\delta)^2}{(1-\delta)^6} + \frac{(1+\delta)^2}{(1-\delta)^4} - 1 \leq \frac{\delta}{(1-\delta)^4} \left(4 \left(\frac{1+\delta}{1-\delta} \right)^2 + 6 \right) \leq \frac{12\delta}{(1-\delta)^4}$$

where the first estimate repeats (129.a), and the second uses that $\left(\frac{1+\delta}{1-\delta}\right)^2 < 1.5$ for $0 \leq \delta \leq 0.1$. Plugging this estimate into (135) results in

$$-\frac{12\delta}{(1-\delta)^4} \mathbf{Q} \preceq \mathbf{S} - \bar{\mathbf{S}} \preceq \frac{12\delta}{(1-\delta)^4} \mathbf{Q}.$$

Joining this with (128), (129), and the bound $\mathbf{Q} \preceq \nabla^2 V_t(w)$, cf. (119), we get (128), as required. \square

Remark C.2. *The constant factor 42 in Lemma C.5 can be improved to 30, exactly recovering the result in [3], through a somewhat more delicate argument following [3], where one separately bounds $\mathbf{Q}_t(\bar{w}) - \mathbf{Q}_t(w)$ and $\mathbf{R}_t(\bar{w}) - \mathbf{R}_t(w)$ for $\mathbf{R}(w) = \nabla^2 V_t(w) - \mathbf{Q}_t(w) = 2\mathbf{Q}_t(w) - 2\mathbf{S}_t(w)$. Here we avoided this, as it would burden the proof with linear algebra and obscure the high-level picture.*

D Stability lemma for Algorithm 1

We are now about to prove Lemma 4.1 about the stability of approximation of the updates according to (VB-FTRL) by those corresponding to (VB-FTRL-qN). For convenience, let us repeat it here.

Lemma D.1. *Let λ, μ, S be as in the premise of Theorem 4.1. Then for all $t \in \{0\} \cup [T]$, one has*

$$\|\tilde{w}_t - w_t\|_{\nabla^2 L_t(w_t)} \leq \min \left\{ \frac{1}{(T+d+1)^2}, 10^{-4}(1+3\mu)^{-1/2} \right\}. \quad (137)$$

Proof. Let $v_t := A^+(w_t - e_d)$, $\tilde{v}_t := A^+(\tilde{w}_t - e_d)$, and $v_t^{(s)} := A^+(w_t^{(s)} - e_d)$ for all $t \in [T]$ and $s \in [S]$. We proceed by induction over $t \in [T]$. The base is obvious: $\tilde{w}_0 = w_0$. For the induction step, assume that (78) holds for some $0 \leq t \leq T-1$. By Proposition C.1 and Lemma A.6, $\sqrt{21}P_t^{\mathbb{1}}(\cdot)$ is a 1-self-concordant function minimized at v_{t+1} . On the other hand, since function $\underline{P}_t^{\mathbb{1}}(\cdot)$ defined in (54) satisfies $\nabla \underline{P}_t^{\mathbb{1}}(v_t) = \nabla P_t^{\mathbb{1}}(v_t)$ and $\nabla^2 \underline{P}_t^{\mathbb{1}}(v_t) \preceq \nabla^2 P_t^{\mathbb{1}}(v_t)$, we estimate

$$\text{Decr}_t := \|\nabla P_t^{\mathbb{1}}(v_t)\|_{\nabla^2 P_t^{\mathbb{1}}(v_t)^{-1}}$$

in the same way in the proof of Theorem 2.1, except for using a stronger condition on λ, μ in (77):

$$\sqrt{21}\text{Decr}_t^2 \leq \sqrt{21}\underline{\text{Decr}}_t^2 \stackrel{(73)}{\leq} \frac{\sqrt{21}}{\lambda} \left(1 + \frac{2\mu}{\lambda} \right)^2 \stackrel{(77)}{\leq} \frac{\sqrt{21}}{556} < \frac{1}{121}. \quad (138)$$

Since $\sqrt{21}\text{Decr}_t^2$ is the squared Newton decrement of $\sqrt{21}P_t^1$ at v_t , by Lemmas A.7–A.8 we get

$$\omega \left(21^{1/4} \|w_t - w_{t+1}\|_{\nabla^2 P_t(w_{t+1})} \right) \stackrel{(11)}{=} \omega \left(21^{1/4} \|v_t - v_{t+1}\|_{\nabla^2 P_t^1(v_{t+1})} \right) \leq \psi \left(21^{1/4} \text{Decr}_t \right)$$

in terms of ω, ψ defined in (83). By monotonicity of these two functions, using (138) we arrive at

$$21^{1/4} \|w_t - w_{t+1}\|_{\nabla^2 P_t(w_{t+1})} \leq \omega^{-1} \left(\psi \left(\frac{1}{11} \right) \right) < 0.097,$$

and then by Lemma A.9 (very conservatively):

$$\nabla^2 P_t(w_{t+1}) \preceq 4\nabla^2 P_t(w_t).$$

Now, observe that $\nabla^2 P_t(w) \preceq \nabla^2 L_t(w) + 3\mu Q_t(w) \preceq (1 + 3\mu)\nabla^2 L_t(w)$, the last step being due to $\lambda \geq 1$ and $\pi_\tau(w_\tau) \leq 1$ by Lemma C.1. Combining this with the previous two inequalities gives

$$\begin{aligned} 21^{1/4} \|\tilde{w}_t - w_{t+1}\|_{\nabla^2 P_t(w_{t+1})} &\leq 21^{1/4} (\|w_t - w_{t+1}\|_{\nabla^2 P_t(w_{t+1})} + 2\|\tilde{w}_t - w_t\|_{\nabla^2 P_t(w_t)}) \\ &< 0.097 + 4.3\sqrt{1 + 3\mu} \|\tilde{w}_t - w_t\|_{\nabla^2 L_t(w_t)} \stackrel{(78)}{<} 0.098. \end{aligned}$$

(In the last step we used the induction hypothesis.) Now, let $\widetilde{\text{Decr}}_t := \|\nabla P_t^1(\tilde{v}_t)\|_{\nabla^2 P_t^1(\tilde{v}_t)^{-1}}$, so that $21^{1/4}\widetilde{\text{Decr}}_t$ is the Newton decrement of $\sqrt{21}P_t^1$ at \tilde{v}_t . By Lemmas A.7–A.8 and (11) we get

$$21^{1/4}\widetilde{\text{Decr}}_t \leq \omega^{-1} \left(\psi \left(21^{1/4} \|\tilde{w}_t - w_{t+1}\|_{\nabla^2 P_t(w_{t+1})} \right) \right) \leq \omega^{-1}(\psi(0.098)) < \frac{1}{9}. \quad (139)$$

We are now in the situation of Corollary A.2 with $f \equiv \sqrt{21}P_t^1$, $z^\star = v_{t+1}$, $z_0 = \tilde{v}_t$, $z_S = \tilde{v}_{t+1}$, and $c = \frac{1}{3}$, cf. (76); in particular, $\|\nabla f(z_0)\|_{\nabla^2 f(z_0)^{-1}} < \frac{c}{3}$ holds by (139). Whence by Corollary A.2:

$$21^{1/4} \|\tilde{w}_{t+1} - w_{t+1}\|_{\nabla^2 P_t(w_{t+1})} \stackrel{(77)}{\leq} \min \left\{ \frac{1}{(T + d + 1)^2}, 10^{-4}(1 + 3\mu)^{-1/2} \right\}.$$

It remains to combine this result with the observation that, due to Lemma A.3, for any $w \in \Delta_d$

$$\nabla^2 L_{t+1}(w) \preceq \nabla^2 L_t(w) + \nabla^2 R(w) \preceq \left(1 + \frac{1}{\lambda}\right) L_t(w) \preceq \left(1 + \frac{1}{\lambda}\right) \nabla^2 P_t(w),$$

and use that $(1 + \frac{1}{\lambda}) \stackrel{(77)}{\leq} (1 + \frac{1}{2e}) < 21^{-1/4}$. □

References

- [1] Amit Agarwal, Elad Hazan, Satyen Kale, and Robert E. Schapire. Algorithms for portfolio management based on the Newton method. In *Proceedings of the 23rd International Conference on Machine learning*, pages 9–16. PMLR, 2006.
- [2] Kurt M. Anstreicher. Large step volumetric potential reduction algorithms for linear programming. *Annals of Operations Research*, 62(1):521–538, 1996.
- [3] Kurt M. Anstreicher. Volumetric path following algorithms for linear programming. *Mathematical Programming*, 76(1):245–263, 1997.

- [4] Ravindra B. Bapat and T. E. S. Raghavan. *Nonnegative matrices and applications*. Cambridge University Press, 1997.
- [5] Pierre C. Bellec. Optimal exponential bounds for aggregation of density estimators. *Bernoulli*, 23(1):219–248, 2017.
- [6] Pierre C. Bellec. Optimal bounds for aggregation of affine estimators. *The Annals of Statistics*, 46(1):30–59, 2018.
- [7] Avrim Blum and Adam Kalai. Universal portfolios with and without transaction costs. *Machine Learning*, 35(3):193–205, 1999.
- [8] Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [9] Nicolas Brosse, Alain Durmus, Éric Moulines, and Marcelo Pereyra. Sampling from a log-concave distribution with compact support with proximal Langevin Monte Carlo. In *Proceedings of the 30th Conference On Learning Theory*, pages 319–342. PMLR, 2017.
- [10] Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Sampling from a log-concave distribution with projected Langevin Monte Carlo. *Discrete & Computational Geometry*, 59(4):757–783, 2018.
- [11] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [12] Thomas M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.
- [13] Thomas M. Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- [14] Thomas M. Cover and Erik Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 1996.
- [15] Arnak S. Dalalyan and Joseph Salmon. Sharp oracle inequalities for aggregation of affine estimators. *The Annals of Statistics*, 40(4):2327–2355, 2012.
- [16] Arnak S. Dalalyan and Alexandre B. Tsybakov. Aggregation by exponential weighting, sharp PAC-Bayesian bounds and sparsity. *Machine Learning*, 72(1):39–61, 2008.
- [17] Raaz Dwivedi, Yuansi Chen, Martin J. Wainwright, and Bin Yu. Log-concave sampling: Metropolis-Hastings algorithms are fast! In *Proceedings of the 31st Conference On Learning Theory*, pages 793–797. PMLR, 2018.
- [18] Gene H. Golub and Charles F. Van Loan. *Matrix computations*. JHU Press, 2013.
- [19] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [20] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2):169–192, 2007.

- [21] David P. Helmbold, Robert E. Schapire, Yoram Singer, and Manfred K. Warmuth. On-line portfolio selection using multiplicative updates. *Mathematical Finance*, 8(4):325–347, 1998.
- [22] Adam T. Kalai and Santosh Vempala. Efficient algorithms for universal portfolios. *Journal of Machine Learning Research*, pages 423–440, 2002.
- [23] John L. Kelly. A new interpretation of information rate. *The Bell System Technical Journal*, 35(4):917–926, 1956.
- [24] Leonid G. Khachiyan. A polynomial algorithm in linear programming. In *Doklady Akademii Nauk*, volume 244, pages 1093–1096. Russian Academy of Sciences, 1979.
- [25] Jyrki Kivinen and Manfred K. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 132(1):1–63, 1997.
- [26] Gilbert Leung and Andrew R. Barron. Information theory and mixing least-squares regressions. *IEEE Transactions on Information Theory*, 52(8):3396–3410, 2006.
- [27] Bin Li and Steven C. H. Hoi. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)*, 46(3):1–36, 2014.
- [28] Bin Li and Steven C. H. Hoi. *Online portfolio selection: Principles and algorithms*. CRC Press, 2018.
- [29] Haipeng Luo, Chen-Yu Wei, and Kai Zheng. Efficient online portfolio with logarithmic regret. In *Proceedings of the 33rd international Conference on Neural Information Processing Systems*, volume 31, pages 8245–8255, 2018.
- [30] Leonard C. MacLean, Edward O. Thorp, and William T. Ziemba. *The Kelly capital growth investment criterion: Theory and practice*, volume 3. World Scientific, 2011.
- [31] Harry M. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.
- [32] Harry M. Markowitz. Foundations of portfolio theory. *The Journal of Finance*, 46(2):469–477, 1991.
- [33] Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Globally convergent newton methods for ill-conditioned generalized self-concordant losses. *Advances in Neural Information Processing Systems*, 32, 2019.
- [34] Zakaria Mhammedi and Alexander Rakhlin. Damped online Newton step for portfolio selection. *arXiv preprint arXiv:2202.07574*, 2022.
- [35] Hariharan Narayanan and Alexander Rakhlin. Efficient sampling from time-varying log-concave distributions. *The Journal of Machine Learning Research*, 18(1):4017–4045, 2017.
- [36] Arkadii S. Nemirovskii and David B. Yudin. Problem complexity and method efficiency in optimization. 1983.
- [37] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.

- [38] Yurii Nesterov and Arkadii S. Nemirovskii. *Interior-point polynomial algorithms in convex programming*. SIAM, 1994.
- [39] Laurent Orseau, Tor Lattimore, and Shane Legg. Soft-Bayes: Prod for mixtures of experts with log-loss. In *International Conference on Algorithmic Learning Theory*, pages 372–399. PMLR, 2017.
- [40] Dmitrii M. Ostrovskii and Francis Bach. Finite-sample analysis of m -estimators using self-concordance. *Electronic Journal of Statistics*, 15(1):326–391, 2021.
- [41] Quentin Paris. Online learning with exponential weights in metric spaces. *arXiv preprint arXiv:2103.14389*, 2021.
- [42] Philippe Rigollet and Alexandre B. Tsybakov. Sparse estimation by exponential weighting. *Statistical Science*, 27(4):558–575, 2012.
- [43] Mark Rubinstein. Markowitz’s “portfolio selection:” a fifty-year retrospective. *The Journal of Finance*, 57(3):1041–1045, 2002.
- [44] Naum Z. Shor and Nikolay G. Zhurbenko. A minimization method using the operation of extension of the space in the direction of the difference of two successive gradients. *Cybernetics*, 7(3):450–459, 1971.
- [45] Pravin M. Vaidya. A new algorithm for minimizing convex functions over convex sets. In *30th Annual Symposium on Foundations of Computer Science*, pages 338–343. IEEE Computer Society, 1989.
- [46] Tim Van Erven, Dirk Van der Hoeven, Wojciech Kotłowski, and Wouter M. Koolen. Open problem: Fast and optimal online portfolio selection. In *Proceedings of the 33rd Conference On Learning Theory*, pages 3864–3869. PMLR, 2020.
- [47] Ramon van Handel. *Probability in High Dimension: Lecture Notes*. Princeton University, 2016.
- [48] Keru Wu, Scott Schmidler, and Yuansi Chen. Minimax mixing time of the Metropolis-adjusted Langevin algorithm for log-concave sampling. *arXiv preprint arXiv:2109.13055*, 2021.
- [49] Julian Zimmert, Naman Agarwal, and Satyen Kale. Pushing the efficiency-regret Pareto frontier for online learning of portfolios and quantum states. *arXiv preprint arXiv:2202.02765*, 2022.
- [50] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, pages 928–936, 2003.