

2021-12-06

Minimax optimal quantile and semi-adversarial regret via root-logarithmic regularizers

J. Negrea, B. Bilodeau, N. Campolongo, F. Orabona, D.M. Roy. 2021. "Minimax Optimal Quantile and Semi-Adversarial Regret via Root-Logarithmic Regularizers." NeurIPS, <https://arxiv.org/abs/2110.14804>
<https://hdl.handle.net/2144/44849>

"Downloaded from OpenBU. Boston University's institutional repository."

Minimax Optimal Quantile and Semi-Adversarial Regret via Root-Logarithmic Regularizers

Jeffrey Negrea*
University of Toronto
jeffrey.negrea@mail.utoronto.ca

Blair Bilodeau*
University of Toronto
blair.bilodeau@mail.utoronto.ca

Nicolò Campolongo
Spanflug Technologies GmbH
nico.campolongo@spanflug.de

Francesco Orabona
Boston University
francesco@orabona.com

Daniel M. Roy
University of Toronto
daniel.roy@utoronto.ca

Abstract

Quantile (and, more generally, KL) regret bounds, such as those achieved by NormalHedge (Chaudhuri, Freund, and Hsu 2009) and its variants, relax the goal of competing against the best individual expert to only competing against a majority of experts on adversarial data. More recently, the semi-adversarial paradigm (Bilodeau, Negrea, and Roy 2020) provides an alternative relaxation of adversarial online learning by considering data that may be neither fully adversarial nor stochastic (i.i.d.). We achieve the minimax optimal regret in both paradigms using FTRL with separate, novel, root-logarithmic regularizers, both of which can be interpreted as yielding variants of NormalHedge. We extend existing KL regret upper bounds, which hold uniformly over target distributions, to possibly uncountable expert classes with arbitrary priors; provide the first full-information lower bounds for quantile regret on finite expert classes (which are tight); and provide an adaptively minimax optimal algorithm for the semi-adversarial paradigm that adapts to the true, unknown constraint faster, leading to uniformly improved regret bounds over existing methods.

1 Introduction

We focus on the setting of learning with expert advice [V90; LW94], where in each round the learner selects a probability distribution over experts, observes the loss of each expert, and incurs the average loss of the experts under the learner’s selected distribution. The learner’s objective is to minimize regret against some mixture of the experts, which is the difference between their cumulative loss and the cumulative loss of the expert mixture over T rounds.

The classical “worst-case” online learning paradigm assumes that the losses are adversarial—that is, they are chosen to make the learner perform as poorly as possible—and demands that the learner competes against the best-performing expert. However, there are many real-world settings where this assumption is too pessimistic, and consequently we focus on designing algorithms with provable guarantees that adapt to easier notions of both data and performance measures. A non-exhaustive list of work on “easy data” includes Freund and Schapire [FS97], Cesa-Bianchi, Mansour, and Stoltz

*Equal contribution authors.

[CMS07], van Erven, Grünwald, Koolen, and de Rooij [v+11], and Gaillard, Stoltz, and van Erven [GSv14], all of which use variants of the Hedge algorithm [FS97] to obtain regret bounds in terms of data-dependent quantities. Ideally, such quantities are small when the data are “easy” to predict.

In this work, we focus on two paradigms beyond the classical worst-case: first, we consider relaxing the performance measure to *quantile (KL) regret*, which measures the ability of an algorithm to compete against an unknown mixture of the experts that potentially performs worse than a point-mass on the single best expert, and second, we consider regret within the *semi-adversarial paradigm*, which defines a spectrum of constraints on the permissible data distributions between stochastic and adversarial. The concept of ϵ -quantile regret was introduced by Chaudhuri, Freund, and Hsu [CFH09], in which the player competes against the $\lfloor \epsilon N \rfloor$ best experts (out of N total) rather than the single best. The authors demonstrated empirically that Hedge does poorly in this paradigm, and introduced a new algorithm NormalHedge with an upper bound on quantile regret of $\sqrt{(T + (\log N)^2)(1 + \log(1/\epsilon))}$. Later algorithms improved it to $\sqrt{T(1 + \log(1/\epsilon))}$ [CV10; OP16], removing the dependency on N . The semi-adversarial paradigm considers constraining the adversary’s choice of data distributions, which was first motivated by Rakhlin, Sridharan, and Tewari [RST11]. Bilodeau, Negrea, and Roy [BNR20] extended this idea, defining *adaptive* minimax regret with respect to such constraints and providing an efficient algorithm with corresponding regret bounds.

Contributions While the best known algorithms for the above two paradigms are intrinsically different, we show that the follow-the-regularized-leader (FTRL) algorithm with new root-logarithmic regularizers achieves minimax optimal performance for both quantile and semi-adversarial regret. First, we provide the first FTRL algorithm with minimax optimal quantile regret guarantees, and do so without using the additive normalization step of previous algorithms [CL06, Section 2.1]. We achieve root-KL bounds that hold uniformly over target distributions on (possibly uncountable) expert classes with arbitrary priors, and reduce to the optimal quantile regret for discrete uniform priors. Moreover, we prove matching lower bounds for quantile regret with finite expert classes, demonstrating the optimality of known upper bounds (including our own) for the first time. Finally, in the semi-adversarial paradigm, we improve the dependence on the number of experts in the regret bound, obtaining uniformly improved upper bounds over previous work.

We achieve the above results through a novel local-norm analysis of FTRL with *linearly decomposable* regularizers on general (possibly uncountable) expert spaces. We use this analysis in conjunction with basic conditions on the first and second derivatives to design and analyze the root-logarithmic regularizers. We believe that this approach is fundamentally different from existing ones and could lead to further advances in obtaining optimal algorithms. In fact, there exist results stating if a regret bound is achievable by some algorithm, then that same bound is nearly achievable by mirror descent with some potential function (see, e.g., Srebro, Sridharan, and Tewari [SST11, Thm 9]). However, it is not clear how to design such a function in practice. In contrast, our general FTRL bound reduces the choice of regularizer to a single univariate function, and clarifies how fundamental properties of this function lead to trade-offs in the regret bound.

Related Work Chernov and Vovk [CV10] first discussed the fact that the ϵ -quantile regret corresponds to the KL divergence between a uniform prior and the uniform mixture of the top $\lfloor \epsilon N \rfloor$ competitors, and such KL bounds have consequently quickly followed quantile regret bounds. Luo and Schapire [LS14; LS15] provided a variant of NormalHedge along with a data-dependent KL regret bound at the cost of an additional $\log T$ factor; a similar but even tighter result was obtained by Koolen and van Erven [Kv15]. Orabona and Pál [OP16] showed that these algorithms can be obtained with a reduction to optimal coin-betting online algorithms, yet none of them can be reduced to an FTRL algorithm. Independently, Harvey, Liaw, Perkins, and Randhawa [H+20] proved that a similar strategy allows one to achieve the optimal anytime regret for the setting with $N = 2$ experts.

The first alternative to additively normalized algorithms in the (bandit) learning with experts setting was the INF algorithm [AB09], which was later recast as online mirror descent [ABL11]. Alquier [A21] also did not use additive normalization, and first introduced FTRL with f -divergences (focusing on the χ^2 -divergence) along with regret bounds for continuous distributions with certain unbounded loss functions. For the specific case of Hedge (and its online mirror descent analogue), an analysis on continuous spaces was given by Krichene, Balandat, Tomlin, and Bayen [K+15] and then more generally by van der Hoeven, van Erven, and Kotłowski [vvK18], while a coarser analysis for f -divergences was given by Alquier [A21]. Finally, choosing regularizers that are tuned to minimize

regret for specific tasks has recently led to advances in the online learning literature, both with bandit feedback [A+17; WL18; F+16] and in the full-information setting [LWZ18].

The best known asymptotic lower bound for learning with expert advice is by Cesa-Bianchi, Freund, Haussler, Helmbold, Schapire, and Warmuth [C+97], while a finite-time lower bound that asymptotically matches the leading constant was proved by Orabona and Pál [OP15a]. We are not aware of lower bounds for the quantile regret, with the notable exception of Koolen [K13], who proved a lower bound on the regret of learning with two experts based on the KL divergence against some prior.

2 Notation

To analyze FTRL beyond the finite setting requires some more care, and we rely on the language of measure theory to handle finite and uncountable expert classes simultaneously. Let (Θ, Σ) be a measurable space, $\mathcal{M}^\infty(\Theta, \Sigma)$ and $\mathcal{M}(\Theta, \Sigma)$ denote the collection of finite and probability measures respectively, and $\nu \in \mathcal{M}^\infty(\Theta, \Sigma)$ be arbitrary. A measure $q \in \mathcal{M}^\infty(\Theta, \Sigma)$ is *absolutely continuous* with respect to ν (denoted $q \ll \nu$) if $q(A) = 0$ for all $A \in \Sigma$ such that $\nu(A) = 0$. Let $\mathcal{M}_\nu(\Theta, \Sigma) = \{q \in \mathcal{M}(\Theta, \Sigma) : q \ll \nu\}$ be the set of probability measures that are absolutely continuous with respect to ν . For any integer N , let $[N] = \{1, \dots, N\}$, and $\text{simp}([N]) = \{p \in [0, 1]^N : \sum_{i=1}^N p_i = 1\}$. Let $\mathbb{R}_+ = [0, \infty)$ and $\mathbb{R}_{++} = (0, \infty)$, and define $\mathbb{M} \equiv \mathbb{M}((\Theta, \Sigma), (\mathbb{R}, \mathcal{R}))$ to be the space of measurable functions from (Θ, Σ) to $(\mathbb{R}, \mathcal{R})$, where \mathcal{R} is the Borel σ -field on \mathbb{R} . Define sets of bounded measurable functions

$$\mathcal{L}^\infty = \{\ell \in \mathbb{M} : \sup_{\theta \in \Theta} |\ell(\theta)| < \infty\} \text{ and } \mathcal{L}_{[0,1]}^\infty = \{\ell \in \mathbb{M} : 0 \leq \ell \leq 1\},$$

sets of integrable functions

$$\mathcal{L}^1(\nu) = \{x \in \mathbb{M} : \int |x(\theta)| \nu(d\theta) < \infty\} \text{ and } \mathcal{L}_+^1(\nu) = \{x \in \mathcal{L}^1(\nu) : x \geq 0\},$$

and the set of Radon–Nikodym derivatives (w.r.t. ν) of probability measures

$$\mathcal{X}(\nu) = \{x \in \mathcal{L}_+^1(\nu) : \int x(\theta) \nu(d\theta) = 1\}.$$

For every $x \in \mathcal{X}(\nu)$, let $\nu^{(x)} \in \mathcal{M}_\nu(\Theta, \Sigma)$ denote the unique probability measure satisfying $d\nu^{(x)}/d\nu = x$. For every $x, z \in \mathcal{L}^\infty$, let $\text{Conv}(x, z) = \{\alpha x + (1 - \alpha)z : \alpha \in [0, 1]\}$ denote their convex hull. For every $\ell \in \mathcal{L}^\infty$ and $x \in \mathcal{L}^1(\nu)$, define

$$\langle \ell, x \rangle_\nu = \int \ell(\theta) x(\theta) \nu(d\theta).$$

Prediction with expert advice We consider the following setting of online linear optimization. For each round $t \in [T]$, the player selects $\mu_t \in \mathcal{M}(\Theta, \Sigma)$ based only upon information available prior to round t , and then observes $\ell_t \in \mathcal{L}^\infty$. Performance is measured using the *regret* against some probability measure $q \in \mathcal{M}(\Theta, \Sigma)$, which is defined by

$$R_T(q) = \sum_{t=1}^T \mathbb{E}_{\theta \sim \mu_t} [\ell_t(\theta)] - \sum_{t=1}^T \mathbb{E}_{\theta \sim q} [\ell_t(\theta)].$$

The player’s goal is to select μ_t so that the cumulative loss is not much larger than that of the average loss of an expert under q , and consequently to have small regret.

The elements of Θ can be regarded as experts for some prediction problem, and the learner aggregates the predictions of the experts by selecting an expert at random according to μ_t at round t . For concreteness, note that the usual prediction with expert advice setting for N experts corresponds to $\mathcal{M}(\Theta, \Sigma) = \text{simp}([N])$. Note that, for any convex loss, applying online linear optimization to the gradients of the losses provides an upper bound on the performance of online convex optimization.

Follow-the-regularized-leader Let $\overline{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$. Follow-the-regularized-leader (FTRL) [S07; AHR08; HK10] forms a broad class of algorithms for online convex optimization. For a finite measure $\nu \in \mathcal{M}^\infty(\Theta, \Sigma)$ and a sequence of regularizers $(\psi_t : \mathcal{X}(\nu) \rightarrow \overline{\mathbb{R}})_{t \in [T]}$, $(\psi_t)_{t \in [T]}$ -regularized FTRL is defined by selecting $\mu_{t+1} = \nu^{(x_{t+1})}$ using

$$x_{t+1} \in \arg \min_{x \in \mathcal{X}(\nu)} \{\langle L_t, x \rangle_\nu + \psi_{t+1}(x)\}, \quad (1)$$

where $L_0 = 0$ and $L_t = \sum_{s=1}^t \ell_s$.

As mentioned above, a classical algorithm for prediction with expert advice on a finite class of N experts is Hedge. While originally analyzed using potential functions, Hedge also corresponds to FTRL where ν is the counting measure and $\psi_t(\cdot) = -\eta_t^{-1}H(\cdot)$ for the Shannon entropy $H(x) = \sum_{i \in [N]} x(i) \log(1/x(i))$ with any sequence of regularizer scalings $(\eta_t)_{t \in \mathbb{N}} \subseteq \mathbb{R}_{++}$.

Depending on whether the space is continuous or discrete, the range of elements of $\mathcal{X}(\nu)$ will change, and consequently also the minimal domain on which regularizers must be defined. Let $\underline{\nu} = \inf \{\nu(A) : A \in \Sigma, \nu(A) > 0\}$, and let $\mathcal{S}(\nu) = [0, 1/\underline{\nu}]$ if $\underline{\nu} > 0$ and \mathbb{R}_+ otherwise. By definition, $\nu(\{\theta \in \Theta : x(\theta) \in \mathcal{S}(\nu)\}) = 1$ for all $x \in \mathcal{X}(\nu)$.

Concretely, when ν is counting measure then $\mathcal{S}(\nu) = [0, 1]$, when ν is uniform on $[N]$ then $\mathcal{S}(\nu) = [0, N]$, and when ν is a continuous distribution then $\mathcal{S}(\nu) = \mathbb{R}_+$. In this work, we consider *linearly decomposable* regularizers of the form $\psi_t(x) = \eta_t^{-1} \Psi_f^\nu(x)$, where

$$\Psi_f^\nu(x) = \int f(x(\theta)) \nu(d\theta) \quad (2)$$

for some $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$. We refer to the algorithm that selects μ_{t+1} using Eq. (1) with a regularizer of the form in Eq. (2) as Ψ_f^ν -regularized FTRL with regularizer scaling η_t .

3 Applications of FTRL with linearly decomposable regularizers

In Section 5, we provide a general analysis of FTRL with linearly decomposable regularizers for arbitrary f . First, we motivate such a general analysis by demonstrating the benefits of studying choices of f beyond traditional FTRL regularizers with multiple examples, including quantile regret. To do so, we state the following corollaries of our general FTRL regret bound (Theorem 2 in Section 5) that achieve “root-KL” and variance bounds respectively. Proofs are deferred to Appendix B.

Corollary 1. *Suppose $\nu \in \mathcal{M}(\Theta, \Sigma)$ and $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ satisfies:*

1. $f(1) = 0$;
2. f is twice continuously differentiable, $f'' > 0$ on \mathbb{R}_{++} , and $f''(0+) > 0$;
3. either f or $1/f''$ is increasing on \mathbb{R}_{++} ;
4. there exist $c_1, c_2 > 0$ such that $f'' \cdot (f + c_1) \geq c_2$.

For any sequence $(\ell_t)_{t \in \mathbb{N}} \subseteq \mathcal{L}_{[0,1]}^\infty$, $q \in \mathcal{M}_\nu(\Theta, \Sigma)$, and T , Ψ_f^ν -regularized FTRL with regularizer scaling $\eta_t = \sqrt{c_2/t}$ achieves

$$R_T(q) \leq \sqrt{\frac{T+1}{c_2}} \Psi_f^\nu(dq/d\nu) + \frac{c_1}{\sqrt{c_2}} \sqrt{T}.$$

The regret bound above can be easily turned into a uniform root-KL bound under the following additional assumption (the result follows from Lemma 3 in Appendix B).

Corollary 2. *If, in addition to the assumptions of Corollary 1, there exist $k_1 \in \mathbb{R}$ and $k_2 \in \mathbb{R}_+$ with $f(x) \leq k_1 + k_2 x \sqrt{\log(1+x)}$, then*

$$R_T(q) \leq \sqrt{\frac{T+1}{c_2}} \left(k_1 + k_2 \sqrt{1 + \text{KL}(q \parallel \pi)} \right) + \frac{c_1}{\sqrt{c_2}} \sqrt{T}.$$

Under different conditions on f , we also derive variance bounds with respect to the (intermediate) predictive distributions, similar to AdaHedge [v+11] and AdaFTRL [OP15b], or with respect to a “prior”, as in [A21]. We state the conditions on the regularizer and the corresponding regret bounds in the following corollary.

Corollary 3. *Suppose $\nu \in \mathcal{M}^\infty(\Theta, \Sigma)$ and $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$ satisfies:*

1. $f(1/\nu(\Theta)) \geq 0$;
2. f is twice continuously differentiable, $f'' > 0$ on $\text{interior}(\mathcal{S}(\nu))$, and $f''(0+) > 0$.

For any sequence $(\ell_t)_{t \in \mathbb{N}} \subseteq \mathcal{L}_{[0,1]}^\infty$, $q \in \mathcal{M}_\nu(\Theta, \Sigma)$, and T , Ψ_f^ν -regularized FTRL achieves

1. If $1/f''(x) \leq Cx$, $\eta_{t+1} = (C[1/2 + \sum_{s=1}^{t-1} \text{Var}_{\theta \sim \nu^{(z_{s+1})}} \ell_s(\theta)])^{-1/2}$ gives

$$R_T(q) \leq [\Psi_f'(dq/d\nu) + 1] \sqrt{C[1/2 + \sum_{t=1}^T \text{Var}_{\theta \sim \nu^{(z_{t+1})}} \ell_t(\theta)]}.$$

2. If $1/f'' \leq C$, $\eta_{t+1} = (C \nu(\Theta)[1/4 + \sum_{s=1}^t \text{Var}_{\theta \sim \bar{\nu}} \ell_s(\theta)])^{-1/2}$ for $\bar{\nu} = \nu/\nu(\Theta)$ gives

$$R_T(q) \leq [\Psi_f'(dq/d\nu) + 1] \sqrt{C \nu(\Theta)[1/4 + \sum_{t=1}^T \text{Var}_{\theta \sim \bar{\nu}} \ell_t(\theta)]}.$$

3.1 Examples of f -divergence FTRL

We now apply all three of these corollaries to more concrete choices of f . Note that in the setting of [Corollary 1](#), since ν is a probability distribution and f is convex with $f(1) = 0$, the regularizer corresponds to an f -divergence, which we denote by $\Psi_f'(dq/d\nu) = D_f(q \parallel \nu)$ for any $q \in \mathcal{M}_\nu(\Theta, \Sigma)$. In this case, we call ν a *prior* and refer to the algorithm as *f -divergence FTRL*. To demonstrate the utility of [Corollary 1](#), we now show how our result recovers the classical analysis for Hedge and applies to our new root-logarithmic regularizer, both examples of f -divergence FTRL. We also use our general expression for the solution to FTRL with linearly decomposable regularizers, given by [Lemma 1](#) in [Section 5](#), to obtain explicit, novel expressions for some solutions of f -divergence FTRL.

Example 1 (Hedge). *The regularizer corresponding to Hedge (and the Gibbs posterior from Bayesian inference) is given by $f(x) = x \log x$, which satisfies $D_f(\cdot \parallel \pi) = \text{KL}(\cdot \parallel \pi)$ for any $\pi \in \mathcal{M}(\Theta, \Sigma)$. This choice satisfies the assumptions of [Corollary 1](#) with $c_1 = 1 + 1/e$ and $c_2 = 1$, since $1/f''$ is increasing, but does not satisfy the conditions of [Corollary 2](#). Thus, we do not obtain uniform root-KL bounds for Hedge; we conjecture this is in fact a limitation of Hedge and not merely an artifact of our analysis. We can, however, apply [Corollary 3](#) with $C = 1$ to obtain*

$$R_T(q) \leq \left[\text{KL}(q \parallel \nu) + 1 \right] \sqrt{1/2 + \sum_{t=1}^T \text{Var}_{\theta \sim \nu^{(z_{t+1})}} \ell_t(\theta)}.$$

Further, by [Lemma 1](#), f -divergence FTRL with regularizer scaling η_{t+1} recovers the familiar formula

$$x_{t+1}(\theta) = \frac{\exp(-\eta_{t+1} L_t(\theta))}{\int \exp(-\eta_{t+1} L_t(\vartheta)) \nu(d\vartheta)}. \quad \triangleleft$$

Example 2 (χ^2 -divergence). *Alquier [A21] analyzed f -divergence FTRL with constant regularizer scalings, obtaining variance bounds ([Theorem 2.1](#)) that require knowledge of the variance to tune the regularizer scaling. He specifically focused on the KL-divergence, covered by the previous example, and the χ^2 -divergence, corresponding to $f(x) = x^2 - 1$. This f clearly satisfies the conditions of [Corollary 1](#) with $c_1 = c_2 = 2$, so we match the optimized bound of Alquier [A21, [Corollary 2.4](#)] with*

$$R_T(q) \leq \sqrt{T} \chi^2(q \parallel \nu) + \sqrt{2T}.$$

Further, the conditions of [Corollary 3](#) are satisfied with $C = 1/2$, so we obtain the novel, potentially much smaller variance bound (without requiring advance knowledge of the variances)

$$R_T(q) \leq \frac{1}{\sqrt{2}} \left[\chi^2(q \parallel \nu) + 1 \right] \sqrt{1/4 + \sum_{t=1}^T \text{Var}_{\theta \sim \nu} \ell_t(\theta)}.$$

By [Lemma 1](#), the χ^2 -divergence FTRL solution with regularizer scaling η_{t+1} is

$$x_{t+1}(\theta) = \frac{1}{2} (k_{t+1}^* - \eta_{t+1} L_t(\theta))_+$$

where $k_{t+1}^* \in \mathbb{R}$ solves $\int [\frac{1}{2}(k_{t+1}^* - \eta_{t+1} L_t(\theta))_+] \nu(d\theta) = 1$, which matches the formula obtained by Alquier [A21, [Example 3.2](#)]. \triangleleft

Example 3 (abNormal). *We call f -divergence FTRL with any f satisfying the conditions of [Corollary 2](#) abNormal. One such example is $f(x) = \int_1^x \sqrt{2 \log(1+s)} ds$, which satisfies the conditions of [Corollary 1](#) with $c_1 = 2$ and $c_2 = 1/\sqrt{2}$ since f is increasing, and satisfies the conditions of [Corollary 2](#) with $k_1 = 0$ and $k_2 = \sqrt{2}$. Thus, we obtain*

$$R_T(q) \leq 2\sqrt{(T+1)(1 + \text{KL}(q \parallel \nu))} + \sqrt{8T}.$$

The f -divergence FTRL solution with regularizer scaling η_{t+1} is given by

$$x_{t+1}(\theta) = \exp \left\{ \left(k_{t+1}^* - \eta_{t+1} L_t(\theta) \right)_+^2 / 2 \right\} - 1,$$

where $k_{t+1}^* \in \mathbb{R}$ solves $\int [\exp \{ (k_{t+1}^* - \eta_{t+1} L_t(\theta))_+^2 / 2 \} - 1] \nu(d\theta) = 1$. Note that this formula is heuristically similar to NormalHedge when $|\Theta| = N$, which assigns weights to experts according to

$$w_t(i) \propto \left[\sum_{s=1}^{t-1} \langle \ell_s, w_s \rangle - L_{t-1}(i) \right]_+ \exp \left(\left[\sum_{s=1}^{t-1} \langle \ell_s, w_s \rangle - L_{t-1}(i) \right]_+^2 / 2c_{t-1} \right),$$

where c_{t-1} solves $\sum_{i \in [N]} \exp \left(\left[\sum_{s=1}^{t-1} \langle \ell_s, w_s \rangle - L_{t-1}(i) \right]_+^2 / 2c_{t-1} \right) = eN$. \triangleleft

3.2 Lower bound for quantile regret

Bounds for quantile regret and KL regret can be related by observing a) that for $\delta_{(i_\epsilon)}$ denoting a point-mass on the i_ϵ th best expert with respect to L_T ,

$$R_T(\delta_{(i_\epsilon)}) \leq R_T(u_\epsilon),$$

where $\epsilon = i_\epsilon/N$ and $u_\epsilon = \frac{1}{i_\epsilon} \sum_{j=1}^{i_\epsilon} \delta_{(j)}$ is the uniform distribution over the top i_ϵ experts, and b) that $\text{KL}(u_\epsilon \parallel \text{Unif}([N])) = \log(1/\epsilon)$. Thus, KL upper bounds are also upper bounds on quantile regret and quantile regret lower bounds are also lower bounds on certain KL regrets. With this in mind, we provide the first general lower bound for quantile regret on N experts. Our lower bound is matching (up to lower order terms) the leading term in our upper bound for quantile regret achieved by abNormal (Example 3) when ν is uniform on $[N]$ and q is uniform on only the top quantile of experts, establishing the minimax rate of quantile regret as $\sqrt{T \log(1/\epsilon)}$.

Theorem 1. For all $N \in \mathbb{N}$ there exists a probability distribution p on $[0, 1]^N$ such that for any sequence of player predictions $(w_t)_{t \in \mathbb{N}} \subseteq \text{simp}([N])$, $i_\epsilon \in \{1, \dots, \lfloor N/4 \rfloor\}$, and $T \in \mathbb{N}$,

$$\mathbb{E}_{\ell_{1:T} \sim p^{\otimes T}} R_T(\delta_{(i_\epsilon)}) \geq \sqrt{(T/2) (\log(1/\epsilon) - 2 \log 2 + 1/\pi)} - \sqrt{2/\pi} - 2 \log N - \log 2,$$

where $\delta_{(i_\epsilon)}$ is the point-mass on the i_ϵ th best expert with respect to L_T and $\epsilon = i_\epsilon/N$.

3.3 Intuition for f -divergence FTRL in the KL regret paradigm

The conditions on f in Corollary 1 are essentially the minimal conditions needed for the summation terms in Theorem 2 to cancel with each other regardless of the actual losses. Intuitively, to achieve f -divergence regret bounds with f -divergence FTRL using the bound of Theorem 2, these terms must cancel so that there is no dependence in the final bound on the regularity of the distributions actually selected by the algorithm. The condition on f in Corollary 2 is essentially the minimal condition needed for Jensen's inequality to imply $D_f(\cdot \parallel \cdot) \lesssim \sqrt{\text{KL}(\cdot \parallel \cdot)}$. That it is possible to satisfy both of these conditions simultaneously is the crucial observation that enables our result.

All regularizers to which both Corollaries 1 and 2 apply are essentially equivalent to $f(x) = \int_1^x \sqrt{2 \log(1+s)} ds$, meaning they have the same asymptotic growth rate. To see this, first observe that this f has the minimum amount of curvature needed to satisfy $f'' \cdot (f + c_1) \geq c_2$ since $(x\sqrt{\log x}) \cdot (x\sqrt{\log x})''$ is asymptotically constant. Second, this f has the largest asymptotic growth rate that still satisfies $f(x) \leq k_1 + k_2 x \sqrt{\log(1+x)}$. These two facts together constrain the shape of f to the root-logarithmic choice.

4 Semi-adversarial regret bounds

We now turn to another perspective on prediction with expert advice for which FTRL with linearly decomposable regularizers is optimal. The semi-adversarial paradigm (introduced by [BNR20]) consists of a family of constraints on the adversary's choice of loss distribution, and the goal of the player is to learn as well as possible for the true constraint without having to know the constraint in advance. More precisely, the setting is characterized by an unknown *time-homogeneous convex constraint* on the adversary's choice of loss distribution, which is formally represented by a convex set of probability distributions on $\mathcal{L}_{[0,1]}^\infty$, denoted by \mathcal{D} . At each round t , the adversary is free to select

any distribution from \mathcal{D} to sample ℓ_t from. Note that when \mathcal{D} is the set of all probability distributions, the worst-case adversarial setting is recovered, and when \mathcal{D} is a singleton, the stochastic setting is recovered. In this section, we describe a new FTRL algorithm (FTRL-CARL) that achieves minimax optimal expected regret without requiring knowledge of \mathcal{D} in advance.

4.1 Adaptive minimax optimality

Minimax regret in the semi-adversarial paradigm is quantified using a few key objects that summarize \mathcal{D} . The first is the collection of *effective stochastic gaps*, defined for each expert $i \in [N]$ by

$$\Delta_i = \inf_{p \in \mathcal{D}} \max_{i' \in [N]} \mathbb{E}_{\ell \sim p} [\ell(i) - \ell(i')].$$

Using these, Bilodeau, Negrea, and Roy [BNR20] define the stochastic gap

$$\Delta_0 = \min \{ \Delta_i : i \in [N], \Delta_i > 0 \}.$$

The second object Bilodeau, Negrea, and Roy [BNR20] define is the set of *effective experts*, which is a subset of $[N]$ defined by

$$\mathcal{I}_0 = \{i \in [N] : \Delta_i = 0\}.$$

This is the set of all experts who are optimal in expectation for some element of \mathcal{D} (or possibly in the limit along some sequence in \mathcal{D}). The number of effective experts is then $N_0 = |\mathcal{I}_0|$.

In this paradigm, the goal is to compete against the best expert. Letting δ_i denote a point-mass on expert i , for notational simplicity we set $R_T \equiv \max_{i \in [N]} R_T(\delta_i)$. Further, although the player does not expect to have knowledge of \mathcal{D} in advance, the goal is to develop methods that do as well as they possibly could have *if they had access to properties of \mathcal{D} in advance*. To characterize this, we say an algorithm (which only has knowledge of N in advance) is *adaptively minimax optimal* if there exists a constant C such that, for all N and (N_0, Δ_0) pairs, the expected regret of the algorithm is within a factor of C from the minimax regret had the algorithm had access to (N_0, Δ_0) in advance, for sufficiently large T (where sufficiently large may depend on N , N_0 , and Δ_0). For a precise mathematical formulation of this concept, see Section 3.1 of Bilodeau, Negrea, and Roy [BNR20].

When $N_0 = 1$, Proposition 4 of Mourtada and Gaïffas [MG19] shows that the minimax regret is of order no smaller than $(\log N)/\Delta_0$, and when $N_0 > 1$, Theorem 2 of Bilodeau, Negrea, and Roy [BNR20] shows that the minimax regret is of order no smaller than $\sqrt{T \log N_0}$. Bilodeau, Negrea, and Roy [BNR20] prove that Hedge is *not* adaptively minimax optimal in the semi-adversarial paradigm, and their argument applies to other similar Hedge-based algorithms, such as prod [CMS07], AdaHedge [v+11], and Adapt-ML-Prod [GSv14]. Further, they provide an algorithm that achieves

$$\mathbb{E}R_T \lesssim \sqrt{T \log N_0} + \mathbb{I}[N_0 = 1] \frac{\log N}{\Delta_0} + \mathbb{I}[N_0 > 1] \frac{(\log N)^{3/2}}{\Delta_0}. \quad (3)$$

Since $\mathbb{I}[N_0 > 1](\log N)^{3/2}\Delta_0^{-1}$ is lower order when $N_0 > 1$, this algorithm is adaptively minimax optimal. We now present a new algorithm, FTRL-CARL, which is also adaptively minimax optimal *and* achieves a better regret bound for small T .

4.2 Semi-adversarial regret bound for FTRL-CARL

Let $h_C : [0, 1] \rightarrow \mathbb{R}$ be given by

$$h_C(x) = \begin{cases} x\sqrt{2\log(1/x)} - \sqrt{\frac{\pi}{2}} \operatorname{erf}\left(\sqrt{\log(1/x)}\right) + x(N-1)\sqrt{\frac{\pi}{2}} & x \in (0, 1] \\ -\sqrt{\pi/2} & x = 0, \end{cases}$$

set $H_C(w) = \sum_{i \in [N]} h_C(w(i))$, and define FTRL-CARL to be FTRL with ν defined as counting measure, regularizer $\Psi_{-h_C}^\nu(w) = -H_C(w)$, and regularizer scaling $\eta_t = 2/\sqrt{t}$. Note that this corresponds to FTRL with a linearly decomposable regularizer, and an intuitive explanation for this choice of regularizer can be found in Subsection 4.3. We then have the following regret bound for FTRL-CARL, which removes the term $\mathbb{I}[N_0 > 1](\log N)^{3/2}\Delta_0^{-1}$ from Eq. (3).

Corollary 4. *For any time-homogeneous convex constraint \mathcal{D} , FTRL-CARL achieves: For all T ,*

$$\mathbb{E}R_T \leq \sqrt{2T \log N},$$

and if $T > 8(\log N_0)\Delta_0^{-2}$,

$$\mathbb{E}R_T \leq \sqrt{2T \log N_0} + 25 \frac{\log N}{\Delta_0}.$$

Corollary 4 follows from Theorem 3 in Appendix D, which is a more refined regret bound.

4.3 Intuition for FTRL-CARL in the semi-adversarial paradigm

The CARL regularizer can be motivated by the following intuition from the Hedge algorithm. For i.i.d. losses, the upper bounds for Hedge [Theorem 2, MG19] are only optimal with regularizer scaling $\eta_t \gtrsim \sqrt{(\log N)/t}$, and matching lower bounds for the adversarial setting suggest this regularizer scaling constraint is actually necessary for optimal performance. Such a regularizer scaling ensures that the weights of each suboptimal (in expectation) expert decay fast enough. In the semi-adversarial paradigm with more than one effective expert, FTRL-CARE (of [BNR20]) can be interpreted as Hedge with an adaptive regularizer scaling that asymptotically satisfies $\eta_t \gtrsim \sqrt{(\log N_0)/t}$. This smaller regularizer scaling applied to the effective experts is necessary to incur asymptotically $\sqrt{T \log N_0}$ regret. However, since this regularizer scaling is smaller, when there are two or more effective experts the weights assigned to the ineffective experts seemingly are slightly too large.

To rectify this, heuristically, it would be ideal to have expert-specific regularizer scalings of size $\sqrt{(\log N_0)/t}$ for the effective experts and $\sqrt{(\log N)/t}$ for the ineffective experts. Since the effective experts will have weights on the order of $1/N_0$ and the ineffective experts will have weights smaller than $1/N$, the expert-specific regularizer scaling $\eta_t(i) = c\sqrt{\log(1/w_t(i))/t}$ may plausibly achieve the desired behaviour.

The weights of Hedge are defined by the equation $\log(1/w_t(i)) = \eta_t(L_{t-1}(i) + \lambda_t)$, where λ_t is chosen to ensure the weights are normalized. Replacing η_t with our heuristic yields $\log(1/w_t(i)) = c\sqrt{\log(1/w_t(i))/t} (L_{t-1}(i) + \lambda_t)$, which can be rearranged to obtain

$$\sqrt{\log(1/w_t(i))} = \frac{c}{\sqrt{t}}(L_{t-1}(i) + \lambda_t).$$

By Lemma 1, this is exactly the formula for the weights produced with FTRL for the regularizer $-H_C$ and regularizer scaling $\eta_t = c/\sqrt{t}$.

As Theorem 3 shows, our modification to the Hedge algorithm is sufficient to yield semi-adversarial regret bounds with expected regret contribution of size $(\log N)/\Delta_0$ from the ineffective experts, improving on the order of the regret bound for FTRL-CARE when $N_0 > 1$.

5 FTRL analysis for general expert spaces

Finally, we analyze the general performance of FTRL with a linearly decomposable regularizer. First, Lemma 1 provides a closed-form expression for the FTRL solution on general spaces. Specifically, it reduces solving the FTRL optimization problem of Eq. (2) to finding a root of a one-dimensional equation (Eq. (4)), rather than solving a complex, possibly infinite dimensional, optimization problem. Second, Theorem 2 provides our fundamental regret bound for FTRL with such regularizers, which we have used to obtain the results in Sections 3 and 4. It generalizes the well-known results involving local norms from the finite dimensional case [AR09; ZS19; O19], and retains an additional summation term (usually uniformly bounded in FTRL analyses) that is crucial both for tight root-KL regret bounds and for tight bounds in the semi-adversarial paradigm. All results in this section are proved in Appendix A. For the remainder of this section, let $\nu \in \mathcal{M}^\infty(\Theta, \Sigma)$ be fixed and arbitrary.

5.1 Computing FTRL with linearly decomposable regularizers

For continuously differentiable $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$, let $m_{f'} = \inf_{x \in \mathcal{S}(\nu)} f'(x)$, $M_{f'} = \sup_{x \in \mathcal{S}(\nu)} f'(x)$, and $\tau_{f'}(y) = \max(\min(y, M_{f'}), m_{f'})$. We focus on the setting where f' is strictly increasing, and thus $\tau_{f'}(y)$ truncates its argument to the domain of $[f']^{-1}$.

Lemma 1. *Suppose $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$ is twice continuously differentiable with $f'' > 0$ on $\text{interior}(\mathcal{S}(\nu))$ and $f''(0+) > 0$. For any $L \in \mathcal{L}^\infty$ and $\eta > 0$,*

$$x^*(\theta) = [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k^*))$$

satisfies $x^* \in \arg \min_{x \in \mathcal{X}(\nu)} \{\langle L, x \rangle_\nu + \eta^{-1} \Psi_f^\nu(x)\}$, where $k^* \in \mathbb{R}$ solves

$$\int [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k^*)) \nu(d\theta) = 1. \quad (4)$$

Further, this solution is unique up to modification on a set of ν -measure 0.

For finite expert classes, solving the normalizing equation provided by Lemma 1 to a given precision is essentially the same difficulty as normalizing the weights for the classical Hedge algorithm. For example, for FTRL-CARL the range of the normalizing constant scales with $\log N$, and solving this using the bisection method would require only $\mathcal{O}(\log \log N)$ times more computation than normalizing Hedge. For uncountable expert classes, many of the general algorithms for prediction with expert advice from previous work are not easily modified to apply, and the standard analyses of regret bounds rely heavily on a finite N . In this setting, solving the FTRL normalizing equation is essentially as difficult as normalizing the posterior distribution for Bayesian inference. That is, often computationally intractable, yet regularly studied for its theoretical properties. Extending approximation techniques for Bayesian inference to approximate the solution to this optimization problem is an interesting question for future work.

5.2 Choosing f to obtain specific regret bounds

In order to state our generic decomposition for the regret of FTRL algorithms on abstract spaces, we rely on the definition of the one-dimensional *Bregman divergence*, which is defined for any continuously differentiable, strictly convex $f : \mathbb{R} \rightarrow \mathbb{R}$ by

$$B_f(x; y) = f(x) - f(y) - f'(y)(x - y).$$

Theorem 2. *Suppose $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$ is twice continuously differentiable with $f'' > 0$ on interior($\mathcal{S}(\nu)$) and $f''(0+) > 0$. For all sequences $(\ell_t)_{t \in \mathbb{N}} \subseteq \mathcal{L}^\infty$ and $(m_t)_{t \in \mathbb{N}} \subseteq \mathbb{R}$, there exist $\hat{z}_{t+1} \in \text{Conv}(x_t, x_{t+1})$ and $\tilde{z}_{t+1} \in \text{Conv}(x_t, \hat{x}_{t+1})$ such that for all $q \in \mathcal{M}_\nu(\Theta, \Sigma)$, $(\tilde{z}_{t+1})_{t \in \mathbb{N}} \in \{(\hat{z}_{t+1})_{t \in \mathbb{N}}, (\tilde{z}_{t+1})_{t \in \mathbb{N}}\}$, and T, Ψ_f^ν -regularized FTRL achieves*

$$\begin{aligned} R_T(q) &\leq \frac{1}{\eta_{T+1}} \Psi_f^\nu \left(\frac{dq}{d\nu} \right) - \frac{1}{\eta_1} \min_{x \in \mathcal{X}(\nu)} \Psi_f^\nu(x) \\ &\quad + \sum_{t=1}^T \left[\int \frac{\eta_t (\ell_t(\theta) - m_t)^2}{2 f''(\tilde{z}_{t+1}(\theta))} \nu(d\theta) - \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \Psi_f^\nu(x_{t+1}) \right], \end{aligned} \quad (5)$$

where

$$\tilde{x}_{t+1} \in \arg \min_{x \in \mathcal{L}_+^1(\nu)} \left\{ \langle \ell_t - m_t \mathbb{1}, x \rangle_\nu + \frac{1}{\eta_t} \int B_f(x(\theta); x_t(\theta)) \nu(d\theta) \right\}.$$

Further, if $\inf_{\theta \in \Theta} \ell_t(\theta) \geq m_t$, then $\tilde{z}_{t+1} \leq x_t$ pointwise.

Theorem 2 is a general result that can be used to prove regret bounds for a variety of settings under the general FTRL framework we describe, as we have already done in Sections 3 and 4. Furthermore, we highlight that the functional form of Theorem 2 makes it clear how to select the regularizers for both the quantile regret and semi-adversarial paradigms. Specifically, in the former, the terms in summation in Eq. (5) are balanced to cancel, while in the latter, the terms in summation in Eq. (5) are balanced to contribute the same order to the regret. These two cases correspond to the relationship

$$f \cdot f'' = \pm 1. \quad (6)$$

Since the generic bound we obtain in Theorem 2 yields components based upon integrals of both f and $1/f''$, balancing these terms without free tuning parameters requires Eq. (6) to approximately hold. Heuristically, trying to balance these terms with tuning parameters rather than by the choice of f seems to lead to non-adaptive or non-uniform bounds. Hedge provides an example of this for both the semi-adversarial and KL regret cases. In the semi-adversarial paradigm, Bilodeau, Negrea, and Roy [BNR20] showed that Hedge cannot be tuned in a way that is minimax optimal and agnostic to the semi-adversarial constraint that prevails. Similarly, without tuning the regularizer scaling to be dependent on the comparator distribution (equivalently, the quantile of interest), KL (and quantile) regret bounds for Hedge are suboptimal.

5.3 Applications of root-KL regret bounds for continuous experts

We now briefly discuss two applications that highlight the immediate benefits of our general analysis (and consequently results of Section 3) applying beyond finite expert spaces.

Predicting as well as the terminal posterior using f -divergence FTRL A reasonable choice of distribution to measure regret against is the posterior distribution $\hat{\pi}_T$ after having seen T rounds of data. A consequence of Corollaries 1 and 2 is that, for bounded log-likelihoods, the total loss incurred by making predictions according to f -divergence FTRL for suitable f is bounded by the loss incurred by the terminal posterior $\hat{\pi}_T$ plus an excess regret of the order $\sqrt{T \text{KL}(\hat{\pi}_T \parallel \pi)}$. This excess loss is smaller than $T + \text{KL}(\hat{\pi}_T \parallel \pi)$, which is the best available bound for excess loss when predicting according to the posterior $\hat{\pi}_t$ at each round $t \in [T]$, and smaller than $\sqrt{T} \text{KL}(\hat{\pi}_T \parallel \pi)$, which is the best available bound for the excess loss when predicting with Hedge if it is not tuned with *a priori* knowledge of $\text{KL}(\hat{\pi}_T \parallel \pi)$. These worse excess loss bounds follow from the analysis contained in Zhang [Z06], although he only compares against the “true” data-generating parameter.

Model selection using f -divergence FTRL Extending the interpretation of Orabona and Pál [OP16] by Foster, Krishnamurthy, and Luo [FKL20] to countable union model classes, we can consider an infinite sequence of disjoint finite expert classes (note any nested sequence can be made disjoint), $\Theta_1, \Theta_2, \dots$, and let $\Theta = \bigcup_{m \geq 1} \Theta_m$. Assigning to Θ the prior $\pi(\theta) \propto \frac{1}{m^2 |\Theta_m|}$ for each $\theta \in \Theta_m$, we recover the following regret bound for f -divergence FTRL with regularizer scaling $\eta_t \propto 1/\sqrt{t}$ from Corollary 2:

$$R_T(\delta_\theta) \leq \mathcal{O}\left(\sqrt{T(\log |\Theta_m| + \log m)}\right) \text{ for all } \theta \in \Theta_m \text{ and all } m \in \mathbb{N}.$$

6 Limitations

One limitation of our results, which is common in the online learning literature, is that they only apply to bounded losses. The extension of learning theory to unbounded losses has seen increased interest in recent years [GM20; A21; MVZ21], although it remains a major open problem to achieve guarantees for arbitrary losses. In particular, log loss for many nonparametric learners is not covered by the current unbounded loss literature, and will have important implications in statistical learning and density estimation when resolved.

A second limitation of our work is that finding the implicit normalizing constant— k^* in Lemma 1—is computationally difficult, as was also observed by Alquier [A21]. It is at least as costly as finding the normalizing constant for a Bayesian posterior (or running Hedge), which corresponds to one evaluation of the left-hand side of Eq. (4). For Hedge, where $f(x) = x \log(x)$, one evaluation of that equation is sufficient, but in general finding the root of Eq. (4) to a fixed precision will take a number of evaluations depending on the desired precision and the range of possible values, and so FTRL with a general f will be that marginally more computationally expensive than Hedge. Variational approaches for Bayesian inference, which avoid direct normalization of the posterior, may also be applicable for FTRL with a general f , and that line of inquiry may lead to novel, efficient, and high-performance learning algorithms.

A final limitation is that we do not obtain variance bounds together with uniform root-KL bounds, although our approach leads to both separately (Corollary 3 and Corollary 2), so we are optimistic our techniques can lead to such bounds, which would resolve multiple open problems.

Acknowledgments and Disclosure of Funding

JN is supported by an NSERC Vanier Canada Graduate Scholarship and the Vector Institute. BB is supported by an NSERC Canada Graduate Scholarship and the Vector Institute. FO is partly supported by the National Science Foundation under grants no. 1908111 "AF: Small: Collaborative Research: New Representations for Learning Algorithms and Secure Computation" and no. 2046096 "CAREER: Parameter-free Optimization Algorithms for Machine Learning". DMR is supported in part by an NSERC Discovery Grant, an Ontario Early Researcher Award, and a stipend provided by the Charles Simonyi Endowment. We thank Mufan Li and Mahdi Haghifam for helpful feedback on early drafts.

References

- [A+17] A. Agarwal, H. Luo, B. Neyshabur, and R. E. Schapire. “Corralling a Band of Bandit Algorithms”. In: *Proceedings of the 30th Conference on Learning Theory*. 2017.
- [A21] P. Alquier. “Non-Exponentially Weighted Aggregation: Regret Bounds for Unbounded Loss Functions”. In: *Proceedings of the 38th International Conference on Machine Learning*. 2021.
- [AB09] J.-Y. Audibert and S. Bubeck. “Minimax Policies for Adversarial and Stochastic Bandits”. In: *Proceedings of the 22nd Conference on Learning Theory*. 2009.
- [ABL11] J.-Y. Audibert, S. Bubeck, and G. Lugosi. “Minimax Policies for Combinatorial Prediction Games”. In: *Proceedings of the 24th Conference on Learning Theory*. 2011.
- [AC65] M. M. Ali and L. K. Chan. “Some Bounds for Expected Values of Order Statistics”. *Annals of Mathematical Statistics* 36.3 (1965), pp. 1055–1057.
- [AHR08] J. D. Abernethy, E. Hazan, and A. Rakhlin. “Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization”. In: *Proceedings of the 21st Conference on Learning Theory*. 2008.
- [AR09] J. D. Abernethy and A. Rakhlin. “Beating the Adaptive Bandit with High Probability”. In: *Proceedings of the 22nd Conference on Learning Theory*. 2009.
- [B42] Z. W. Birnbaum. “An Inequality for Mill’s Ratio”. *Annals of Mathematical Statistics* 13.2 (1942), pp. 245–246.
- [B60] H. Bauer. “Minimalstellen von Funktionen und Extrempunkte. II”. *Archiv der Mathematik* 11.1 (1960), pp. 200–205.
- [BM89] J. Bretagnolle and P. Massart. “Hungarian Constructions from the Nonasymptotic Viewpoint”. *Annals of Probability* 17.1 (1989), pp. 239–256.
- [BNR20] B. Bilodeau, J. Negrea, and D. M. Roy. *Relaxing the I.I.D. Assumption: Adaptively Minimax Optimal Regret via Root-Entropic Regularization*. arXiv:2007.06552. 2020.
- [C+97] N. Cesa-Bianchi, Y. Freund, D. Haussler, D. P. Helmbold, R. E. Schapire, and M. K. Warmuth. “How to Use Expert Advice”. *Journal of the ACM* 44.3 (1997), pp. 427–485.
- [CFH09] K. Chaudhuri, Y. Freund, and D. J. Hsu. “A Parameter-Free Hedging Algorithm”. In: *Advances in Neural Information Processing Systems* 22. 2009.
- [CL06] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge: Cambridge University Press, 2006.
- [CMS07] N. Cesa-Bianchi, Y. Mansour, and G. Stoltz. “Improved Second-Order Bounds for Prediction with Expert Advice”. *Machine Learning* 66.2 (2007), pp. 321–352.
- [CV10] A. Chernov and V. G. Vovk. “Prediction with Advice of Unknown Number of Experts”. In: *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*. 2010.
- [d+14] S. de Rooij, T. van Erven, P. D. Grünwald, and W. M. Koolen. “Follow the Leader If You Can, Hedge If You Must”. *Journal of Machine Learning Research* 15.37 (2014), pp. 1281–1316.
- [F+16] D. J. Foster, Z. Li, T. Lykouris, K. Sridharan, and É. Tardos. “Learning in Games: Robustness of Fast Convergence”. In: *Advances in Neural Information Processing Systems* 30. 2016.
- [FKL20] D. J. Foster, A. Krishnamurthy, and H. Luo. “Open Problem: Model Selection for Contextual Bandits”. In: *Proceedings of the 33rd Conference on Learning Theory*. 2020.
- [FS97] Y. Freund and R. Schapire. “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. *Journal of Computer and System Sciences* 55.1 (1997), pp. 119–139.
- [GM20] P. D. Grünwald and N. A. Mehta. “Fast Rates for General Unbounded Loss Functions: From ERM to Generalized Bayes”. *Journal of Machine Learning Research* 21.56 (2020), pp. 1–80.
- [GSv14] P. Gaillard, G. Stoltz, and T. van Erven. “A Second-Order Bound with Excess Losses”. In: *Proceedings of the 27th Conference on Learning Theory*. 2014.
- [H+20] N. J. A. Harvey, C. Liaw, E. Perkins, and S. Randhawa. “Optimal Anytime Regret with Two Experts”. In: *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science*. 2020.

- [HK10] E. Hazan and S. Kale. “Extracting Certainty from Uncertainty: Regret Bounded by Variation in Costs”. *Machine Learning* 80.2 (2010), pp. 165–188.
- [K+15] W. Krichene, M. Balandat, C. Tomlin, and A. Bayen. “The Hedge Algorithm on a Continuum”. In: *Proceedings of the 32nd International Conference on Machine Learning*. 2015.
- [K13] W. M. Koolen. “The Pareto Regret Frontier”. In: *Advances in Neural Information Processing Systems* 26. 2013.
- [Kv15] W. M. Koolen and T. van Erven. “Second-Order Quantile Methods for Experts and Combinatorial Games”. In: *Proceedings of the 28th Conference on Learning Theory*. 2015.
- [L69] D. G. Luenberger. *Optimization by Vector Space Methods*. Wiley, 1969.
- [LS14] H. Luo and R. E. Schapire. “A Drifting-Games Analysis for Online Learning and Applications to Boosting”. In: *Advances in Neural Information Processing Systems* 27. 2014.
- [LS15] H. Luo and R. E. Schapire. “Achieving All with no Parameters: Adanormalhedge”. In: *Proceedings of the 28th Conference on Learning Theory*. 2015.
- [LW94] N. Littlestone and M. K. Warmuth. “The Weighted Majority Algorithm”. *Information and Computation* 108.2 (1994), pp. 212–261.
- [LWZ18] H. Luo, C.-Y. Wei, and K. Zheng. “Efficient Online Portfolio with Logarithmic Regret”. In: *Advances in Neural Information Processing Systems* 32. 2018.
- [MG19] J. Mourtada and S. Gaïffas. “On the Optimality of the Hedge Algorithm in the Stochastic Regime.” *Journal of Machine Learning Research* 20.83 (2019), pp. 1–28.
- [MVZ21] J. Mourtada, T. Vaškevičius, and N. Zhivotovskiy. *Distribution-Free Robust Linear Regression*. arXiv:2102.12919. 2021.
- [O19] F. Orabona. *A Modern Introduction to Online Learning*. arXiv:1912.13213. 2019.
- [OP15a] F. Orabona and D. Pál. *Optimal Non-Asymptotic Lower Bound on the Minimax Regret of Learning with Expert Advice*. arXiv:1511.02176. 2015.
- [OP15b] F. Orabona and D. Pál. “Scale-Free Algorithms for Online Linear Optimization”. In: *Proceedings of the 26th International Conference on Algorithmic Learning Theory*. 2015.
- [OP16] F. Orabona and D. Pál. “Coin Betting and Parameter-Free Online Learning”. In: *Advances in Neural Information Processing Systems* 29. 2016.
- [RST11] A. Rakhlin, K. Sridharan, and A. Tewari. “Online Learning: Stochastic, Constrained, and Smoothed adversaries”. In: *Advances in Neural Information Processing Systems* 25. 2011.
- [S07] S. Shalev-Shwartz. “Online Learning: Theory, Algorithms, and Applications”. PhD thesis. The Hebrew University, 2007.
- [SST11] N. Srebro, K. Sridharan, and A. Tewari. “On the Universality of Online Mirror Descent”. In: *Advances in Neural Information Processing Systems* 24. 2011.
- [v+11] T. van Erven, P. D. Grünwald, W. M. Koolen, and S. de Rooij. “Adaptive Hedge”. In: *Advances in Neural Information Processing Systems* 25. 2011.
- [V90] V. G. Vovk. “Aggregating Strategies”. In: *Proceedings of 3rd Workshop on Computational Learning Theory*. 1990.
- [vvK18] D. van der Hoeven, T. van Erven, and W. Kotłowski. “The Many Faces of Exponential Weights in Online Learning”. In: *Proceedings of the 31st Conference On Learning Theory*. 2018.
- [WL18] C.-Y. Wei and H. Luo. “More Adaptive Algorithms for Adversarial Bandits”. In: *Proceedings of the 31st Conference On Learning Theory*. 2018.
- [Z06] T. Zhang. “Information-Theoretic Upper and Lower Bounds for Statistical Estimation”. *IEEE Transactions on Information Theory* 52.4 (2006), pp. 1307–1321.
- [ZS19] J. Zimmert and Y. Seldin. “An Optimal Algorithm for Stochastic and Adversarial Bandits”. In: *The 22nd International Conference on Artificial Intelligence and Statistics*. 2019.

A Proofs for FTRL analysis for general expert spaces

We first prove the formula for the FTRL solution prescribed by [Lemma 1](#), which we recall next.

Lemma 1. *Suppose $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$ is twice continuously differentiable with $f'' > 0$ on $\text{interior}(\mathcal{S}(\nu))$ and $f''(0+) > 0$. For any $L \in \mathcal{L}^\infty$ and $\eta > 0$,*

$$x^*(\theta) = [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k^*))$$

satisfies $x^ \in \arg \min_{x \in \mathcal{X}(\nu)} \{\langle L, x \rangle_\nu + \eta^{-1} \Psi_f^\nu(x)\}$, where $k^* \in \mathbb{R}$ solves*

$$\int [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k^*)) \nu(d\theta) = 1. \quad (4)$$

Further, this solution is unique up to modification on a set of ν -measure 0.

In the proof below, we treat the formula for the optimal solution as an educated guess (inspired by the finite-dimensional case using Lagrange multipliers). The proof then verifies that the proposed solution is well-defined and in fact achieves the optimum. It is also possible to derive the optimal solution directly using a version of the Lagrange multiplier method for convex optimization problems on Banach spaces with cone constraints (see, for example, [\[L69\]](#)). However, to apply such a result, one must still verify the existence of Lagrange multipliers. For example, existence of the Lagrange multiplier for the constraint “ x integrates to 1” exactly corresponds to existence of a solution to [Eq. \(4\)](#). Thus it is not significantly more or less laborious to take such an approach over the “guess and check” method we have employed.

Proof of Lemma 1. Let $\underline{L} = \inf_{\theta \in \Theta} L(\theta)$ and $\bar{L} = \sup_{\theta \in \Theta} L(\theta)$. Since $f'' > 0$, f' is strictly increasing and restricted to a non-negative domain, so $[f']^{-1}$ exists, is strictly increasing, and is non-negative. Let $a = f'(1/\nu(\Theta))$, and note that since $1/\nu(\Theta) \in \mathcal{S}(\nu)$, $\tau_{f'}(a) = a$. Let $g : \mathbb{R} \rightarrow \nu(\Theta) \cdot \mathcal{S}(\nu)$ be given by

$$g(k) = \int [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k)) \nu(d\theta).$$

First, note that for all k , $g(k) \leq \nu(\Theta) \sup_{\theta \in \Theta} [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k))$. So, since $L \geq \underline{L}$, $g(k) \leq \nu(\Theta) [f']^{-1}(\tau_{f'}(-\eta \underline{L} + k))$. Thus,

$$\begin{aligned} 1 &= \int [f']^{-1}(a) \nu(d\theta) \\ &\leq \int [f']^{-1}(\tau_{f'}(-\eta L(\theta) + \eta \bar{L} + a)) \nu(d\theta) \\ &= g(\eta \bar{L} + a) \\ &\leq \nu(\Theta) [f']^{-1}(\tau_{f'}(\eta(\bar{L} - \underline{L}) + a)), \end{aligned} \quad (7)$$

where the second inequality follows since $L \leq \bar{L}$. Further, non-negativity of $[f']^{-1}$ gives

$$\begin{aligned} 0 &\leq g(\eta \underline{L} + a) \\ &= \int [f']^{-1}(\tau_{f'}(\eta(\underline{L} - L(\theta)) + a)) \nu(d\theta) \\ &\leq \int [f']^{-1}(a) \nu(d\theta) \\ &= 1. \end{aligned} \quad (8)$$

By Leibniz rule and the inverse function theorem, [Eqs. \(7\) and \(8\)](#) imply that, for all $k \in [\eta \underline{L} + a, \eta \bar{L} + a]$,

$$\begin{aligned} g'(k) &= \int \frac{\mathbb{I}[M_{f'} \geq -\eta L(\theta) + k \geq m_{f'}]}{f'' \circ [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k))} \nu(d\theta) \\ &\leq \frac{\nu(\Theta)}{\inf_{\theta \in (0, [f']^{-1}(\tau_{f'}(\eta(\bar{L} - \underline{L}) + a))} f''(\theta)} \\ &< \infty. \end{aligned}$$

Hence, g is (Lipschitz) continuous on this interval, and then, by the intermediate value theorem, there exists k^* solving Eq. (4).

Let $F : \mathcal{L}_+^1(\nu) \rightarrow \mathbb{R}$ be given by

$$F(x) = \langle L, x \rangle_\nu + \frac{1}{\eta} \Psi_f'(x).$$

To avoid notational clutter, for the remainder of the proof all statements involving θ implicitly hold only ν -a.s. Let $x(\theta) = [f']^{-1}(\tau_{f'}(-\eta L(\theta) + k^*))$ and note that the definition of k^* implies that $\int x(\theta) \nu(d\theta) = 1$. We now argue that $-\eta L + k^* \leq M_{f'}$.

If $\underline{\nu} = 0$, then $M_{f'} = \infty$, so trivially $-\eta L + k^* \leq M_{f'}$.

If $\underline{\nu} \neq 0$, then ν is purely atomic and, since f' is increasing, $M_{f'} = f'(1/\underline{\nu})$. By definition of $\tau_{f'}$, $x \leq 1/\underline{\nu}$. If $x < 1/\underline{\nu}$, then $\tau_{f'}(-\eta L(\theta) + k^*) < M_{f'}$, so $-\eta L(\theta) + k^* \leq M_{f'}$. Otherwise, there exists an atom $A \in \Sigma$ for ν such that $x(\theta) = 1/\underline{\nu}$ for $\theta \in A$. Since $x \in \mathcal{X}(\nu)$, we must have $x(\theta) = \mathbb{I}[\theta \in A]/\underline{\nu}$; equivalently x is the Radon–Nikodym derivative of a single-atom probability measure completely concentrated on A . Without loss of generality, $A = \{\theta : x(\theta) = 1/\underline{\nu}\}$ and

$$f'(x(\theta)) = m_{f'} + (M_{f'} - m_{f'}) \mathbb{I}[\theta \in A].$$

Then, any k^* that satisfies

$$\begin{cases} -\eta L(\theta) + k^* \geq M_{f'}, & \theta \in A \\ -\eta L(\theta) + k^* \leq m_{f'}, & \theta \notin A \end{cases}$$

is valid and gives the same solution. In particular, since there is some k^* satisfying this, then $\tilde{k}^* = M_{f'} + \eta[\nu\text{-ess sup}]_{\theta \in A} L(\theta) \leq k^* \leq m_{f'} + \eta[\nu\text{-ess inf}]_{\theta \notin A} L(\theta)$, which implies k^* is a valid solution. Finally, we claim that L must be constant on A . To see this, suppose otherwise there is some \tilde{L} such that for $\theta \in \tilde{A} \subseteq A$, $L(\theta) < \tilde{L}$, and for $\theta \in A \setminus \tilde{A}$, $L(\theta) \geq \tilde{L}$. Then the inverse images of (∞, \tilde{L}) and $[\tilde{L}, \infty)$ must partition A since L is measurable, but one of them must have ν measure zero since A is an atom. Thus, it is also true that $\tilde{k}^* = M_{f'} + \eta[\nu\text{-ess inf}]_{\theta \in A} L(\theta)$, so we have argued that, without loss of generality, in all cases $-\eta L + k^* \leq M_{f'}$.

With this in mind, consider any $z \in \mathcal{X}(\nu)$. Since f is continuously differentiable and convex,

$$\begin{aligned} F(z) &= F(x) + \langle L, z - x \rangle_\nu + \frac{1}{\eta} \int [f(z(\theta)) - f(x(\theta))] \nu(d\theta) \\ &\geq F(x) + \langle L, z - x \rangle_\nu + \frac{1}{\eta} \int f'(x(\theta)) [z(\theta) - x(\theta)] \nu(d\theta) \\ &= F(x) + \int_{x(\theta) > 0} L(\theta) [z(\theta) - x(\theta)] \nu(d\theta) + \int_{x(\theta) = 0} L(\theta) [z(\theta) - x(\theta)] \nu(d\theta) \\ &\quad + \frac{1}{\eta} \int_{x(\theta) > 0} f'(x(\theta)) [z(\theta) - x(\theta)] \nu(d\theta) + \frac{1}{\eta} \int_{x(\theta) = 0} f'(x(\theta)) [z(\theta) - x(\theta)] \nu(d\theta) \\ &= F(x) + \int_{x(\theta) > 0} L(\theta) [z(\theta) - x(\theta)] \nu(d\theta) + \int_{x(\theta) = 0} L(\theta) [z(\theta) - x(\theta)] \nu(d\theta) \\ &\quad + \frac{1}{\eta} \int_{x(\theta) > 0} (-\eta L(\theta) + k^*) [z(\theta) - x(\theta)] \nu(d\theta) + \frac{1}{\eta} \int_{x(\theta) = 0} m_{f'} [z(\theta) - x(\theta)] \nu(d\theta) \\ &= F(x) + \int_{x(\theta) > 0} \frac{k^*}{\eta} [z(\theta) - x(\theta)] \nu(d\theta) + \int_{x(\theta) = 0} \left(L(\theta) + \frac{m_{f'}}{\eta} \right) [z(\theta) - x(\theta)] \nu(d\theta) \\ &\geq F(x) + \int_{x(\theta) > 0} \frac{k^*}{\eta} [z(\theta) - x(\theta)] \nu(d\theta) + \int_{x(\theta) = 0} \frac{k^*}{\eta} [z(\theta) - x(\theta)] \nu(d\theta) \\ &= F(x), \end{aligned} \tag{9}$$

where in the second last step we have used that if $x(\theta) = 0$ then $-\eta L(\theta) + k^* \leq m_{f'}$ and $z(\theta) - x(\theta) \geq 0$. Thus, x is a solution to the FTRL equation. Further, since f is strictly convex, equality can only hold in Eq. (9) if $\nu(x = z) = 1$. \square

To prove [Theorem 2](#), we need the analogue of the finite-dimensional first-order optimality condition. For any $V \subseteq \mathcal{L}^1(\nu)$ and any $F : V \rightarrow \mathbb{R}$ and $x \in \mathcal{L}^1(\nu)$, the Gateaux derivative (in the direction of $z \in \mathcal{L}^1(\nu)$) is

$$\delta F[x; z] = \lim_{\alpha \rightarrow 0} \frac{F(x + \alpha z) - F(x)}{\alpha}.$$

The following result is straightforward: we include a proof for completeness.

Lemma 2. *If $V \subseteq \mathcal{L}^1(\nu)$ is convex and $x = \arg \min_{z \in V} F(z)$, then $\delta F[x; z - x] \geq 0$ for all $z \in V$ where the limit exists.*

Proof of Lemma 2. Towards a contradiction, suppose there exists $z \in V$ with $\delta F[x; z - x] < 0$. Define $z_\alpha = \alpha z + (1 - \alpha)x$ for all $\alpha \in [0, 1]$. By definition,

$$\delta F[x; z - x] = \lim_{\alpha \rightarrow 0} \frac{F(x + \alpha(z - x)) - F(x)}{\alpha} = \lim_{\alpha \rightarrow 0} \frac{F(z_\alpha) - F(x)}{\alpha}.$$

By assumption, this implies that for some $\alpha > 0$, $F(z_\alpha) < F(x)$. However, since $z_\alpha \in V$ for all $\alpha \in [0, 1]$ by convexity, this contradicts the optimality of x . \square

Next, define the functional Bregman divergence for any convex $V \subseteq \mathcal{L}^1(\nu)$ and any $F : V \rightarrow \mathbb{R}$ and $x, z \in \mathcal{L}^1(\nu)$ by

$$B_F(x; z) = F(x) - F(z) - \delta F[z; x - z].$$

We are now able to prove our generic FTRL bound, [Theorem 2](#), which we recall here for completeness.

Theorem 2. *Suppose $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$ is twice continuously differentiable with $f'' > 0$ on $\text{interior}(\mathcal{S}(\nu))$ and $f''(0+) > 0$. For all sequences $(\ell_t)_{t \in \mathbb{N}} \subseteq \mathcal{L}^\infty$ and $(m_t)_{t \in \mathbb{N}} \subseteq \mathbb{R}$, there exist $\hat{z}_{t+1} \in \text{Conv}(x_t, x_{t+1})$ and $\tilde{z}_{t+1} \in \text{Conv}(x_t, \hat{x}_{t+1})$ such that for all $q \in \mathcal{M}_\nu(\Theta, \Sigma)$, $(\bar{z}_{t+1})_{t \in \mathbb{N}} \in \{(\hat{z}_{t+1})_{t \in \mathbb{N}}, (\tilde{z}_{t+1})_{t \in \mathbb{N}}\}$, and T , Ψ_f^ν -regularized FTRL achieves*

$$\begin{aligned} R_T(q) &\leq \frac{1}{\eta_{T+1}} \Psi_f^\nu \left(\frac{dq}{d\nu} \right) - \frac{1}{\eta_1} \min_{x \in \mathcal{X}(\nu)} \Psi_f^\nu(x) \\ &\quad + \sum_{t=1}^T \left[\int \frac{\eta_t}{2} \frac{(\ell_t(\theta) - m_t)^2}{f''(\bar{z}_{t+1}(\theta))} \nu(d\theta) - \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \Psi_f^\nu(x_{t+1}) \right], \end{aligned} \tag{5}$$

where

$$\tilde{x}_{t+1} \in \arg \min_{x \in \mathcal{L}_+^1(\nu)} \left\{ \langle \ell_t - m_t \mathbb{1}, x \rangle_\nu + \frac{1}{\eta_t} \int B_f(x(\theta); x_t(\theta)) \nu(d\theta) \right\}.$$

Further, if $\inf_{\theta \in \Theta} \ell_t(\theta) \geq m_t$, then $\tilde{z}_{t+1} \leq x_t$ pointwise.

Proof of Theorem 2. For all t , let

$$F_t(x) = \langle L_{t-1}, x \rangle_\nu + \frac{1}{\eta_t} \Psi_f^\nu(x).$$

By the trivial extension of Orabona [[O19](#), Lemma 7.1] beyond \mathbb{R}^d , we have

$$R_T(u) \leq \frac{1}{\eta_{T+1}} \Psi_f^\nu \left(\frac{dq}{d\nu} \right) - \frac{1}{\eta_1} \min_{x \in \mathcal{X}(\nu)} \Psi_f^\nu(x) + \sum_{t=1}^T [F_t(x_t) - F_{t+1}(x_{t+1}) + \langle \ell_t, x_t \rangle].$$

Then,

$$\begin{aligned} &F_t(x_t) - F_{t+1}(x_{t+1}) + \langle \ell_t, x_t \rangle \\ &= F_t(x_t) - F_t(x_{t+1}) + \langle \ell_t, x_t - x_{t+1} \rangle - \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \Psi_f^\nu(x_{t+1}). \end{aligned} \tag{10}$$

By [Lemma 2](#) and the definition of x_t ,

$$B_{F_t}(x_{t+1}; x_t) \leq F_{t+1}(x_{t+1}) - F_t(x_t).$$

Thus, substituting this into Eq. (10),

$$\begin{aligned} & F_t(x_t) - F_{t+1}(x_{t+1}) + \langle \ell_t, x_t \rangle \\ & \leq -\frac{1}{\eta_t} B_{\Psi_f^\nu(\cdot)}(x_{t+1}; x_t) + \langle \ell_t, x_t - x_{t+1} \rangle - \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \Psi_f^\nu(x_{t+1}), \end{aligned}$$

where we have used linearity of the functional Bregman divergence.

Next, for any $x, z \in \mathcal{X}(\nu)$,

$$\delta \Psi_f^\nu(\cdot)[z; x - z] = \lim_{\alpha \rightarrow 0} \int \frac{[f(z(\theta) + \alpha(x(\theta) - z(\theta))) - f(z(\theta))]}{\alpha} \nu(d\theta). \quad (11)$$

For any $\theta \in \Theta$ and $\alpha \in (0, 1)$, since f' is increasing, by the mean value theorem

$$\begin{aligned} & f' \left(\min \left\{ \inf_{\theta' \in \Theta} z(\theta') + \alpha(x(\theta') - z(\theta')), \inf_{\theta' \in \Theta} z(\theta') \right\} \right) \\ & \leq \frac{[f(z(\theta) + \alpha(x(\theta) - z(\theta))) - f(z(\theta))]}{\alpha} \\ & \leq f' \left(\max \left\{ \sup_{\theta' \in \Theta} z(\theta') + \alpha(x(\theta') - z(\theta')), \sup_{\theta' \in \Theta} z(\theta') \right\} \right). \end{aligned} \quad (12)$$

Recall the form of x_t given by Lemma 1. For each round t , $\sup_{\theta \in \Theta} \{-\eta_t L_{t-1}(\theta) + k^*\} < \infty$ so $\sup_{\theta \in \Theta} \max \{x_t(\theta), x_{t+1}(\theta)\} < \infty$. Thus, since by continuity f' is finite on $\mathcal{S}(\nu)$, for $x = x_{t+1}$ and $z = x_t$ the upper bound of Eq. (12) is finite. If $f'(0+) > -\infty$, the lower bound is also finite. If $f'(0+) = -\infty$, then since L_{t-1} is bounded, x_t is bounded away from 0 uniformly and thus the lower bound of Eq. (12) is still finite. Thus, for any f satisfying the conditions of the lemma, we can apply the bounded convergence theorem to Eq. (11) to obtain

$$B_{\Psi_f^\nu(\cdot)}(x_{t+1}; x_t) = \int B_f(x_{t+1}(\theta); x_t(\theta)) \nu(d\theta).$$

Thus,

$$\begin{aligned} & F_t(x_t) - F_{t+1}(x_{t+1}) + \langle \ell_t, x_t \rangle_\nu \\ & \leq \int (\ell_t(\theta) - m_t) (x_t(\theta) - x_{t+1}(\theta)) - \frac{1}{\eta_t} B_f(x_{t+1}(\theta); x_t(\theta)) \\ & \quad - \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) f(x_{t+1}(\theta)) \nu(d\theta), \end{aligned}$$

where we have used that any m_t can be added since $x_t \in \mathcal{X}(\nu)$ for all t .

We split the proof now to consider the two choices of intermediate points.

Intermediate Point A

By the mean value remainder form of Taylor's theorem there exists $\hat{z}_{t+1}(\theta) \in \text{Conv}(x_t(\theta), x_{t+1}(\theta))$ such that

$$B_f(x_{t+1}(\theta); x_t(\theta)) = \frac{1}{2} f''(\hat{z}_{t+1}(\theta)) (x_t(\theta) - x_{t+1}(\theta))^2.$$

That is,

$$\begin{aligned} & F_t(x_t) - F_{t+1}(x_{t+1}) + \langle \ell_t, x_t \rangle_\nu \\ & \leq \int (\ell_t(\theta) - m_t) (x_t(\theta) - x_{t+1}(\theta)) - \frac{1}{2\eta_t} f''(\hat{z}_{t+1}(\theta)) (x_t(\theta) - x_{t+1}(\theta))^2 \\ & \quad - \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) f(x_{t+1}(\theta)) \nu(d\theta). \end{aligned} \quad (13)$$

Since f'' is bounded away from zero, we can apply Fenchel-Young point-wise for each $\theta \in \Theta$ to obtain

$$(\ell_t(\theta) - m_t) (x_t(\theta) - x_{t+1}(\theta)) \leq \frac{\eta_t (\ell_t(\theta) - m_t)^2}{2 f''(\hat{z}_{t+1}(\theta))} + \frac{1}{2\eta_t} f''(\hat{z}_{t+1}(\theta)) (x_t(\theta) - x_{t+1}(\theta))^2.$$

Substituting into Eq. (13) and summing over t gives the result.

Intermediate Point B
Alternatively, let

$$\tilde{x}_{t+1} \in \arg \max_{x \in \mathcal{L}_+^1(\nu)} \left\{ -\langle \ell_t - m_t \mathbb{1}, \nu^{(x)} \rangle - \frac{1}{\eta_t} \int B_f(x(\theta); x_t(\theta)) \nu(d\theta) \right\},$$

so

$$\begin{aligned} & F_t(x_t) - F_{t+1}(x_{t+1}) + \langle \ell_t, x_t \rangle_\nu \\ & \leq \int (\ell_t(\theta) - m_t) (x_t(\theta) - \tilde{x}_{t+1}(\theta)) \\ & \quad - \frac{1}{\eta_t} B_f(\tilde{x}_{t+1}(\theta); x_t(\theta)) - \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) f(x_{t+1}(\theta)) \nu(d\theta). \end{aligned}$$

Again by the mean value remainder form of Taylor's theorem, there exists $\tilde{z}_{t+1}(\theta) \in \text{Conv}(x_t(\theta), \tilde{x}_{t+1}(\theta))$ such that

$$B_f(\tilde{x}_{t+1}(\theta); x_t(\theta)) = \frac{1}{2} f''(\tilde{z}_{t+1}(\theta)) (x_t(\theta) - \tilde{x}_{t+1}(\theta))^2.$$

The statement follows from the same point-wise application of Fenchel-Young as for intermediate point A.

Finally, note that for each t and $\theta \in \Theta$, solving the gradient equation and applying a very similar argument to the proof of Lemma 1 gives

$$\tilde{x}_{t+1}(\theta) = [f']^{-1}(\tau_{f'}(-\eta_t(\ell_t(\theta) - m_t) + f'(x_t(\theta)))) \leq x_t(\theta),$$

as long as $\inf_{\theta \in \Theta} \ell_t(\theta) \geq m_t$. □

B Proofs for f -divergence FTRL bounds

Corollary 1. Suppose $\nu \in \mathcal{M}(\Theta, \Sigma)$ and $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ satisfies:

1. $f(1) = 0$;
2. f is twice continuously differentiable, $f'' > 0$ on \mathbb{R}_{++} , and $f''(0+) > 0$;
3. either f or $1/f''$ is increasing on \mathbb{R}_{++} ;
4. there exist $c_1, c_2 > 0$ such that $f'' \cdot (f + c_1) \geq c_2$.

For any sequence $(\ell_t)_{t \in \mathbb{N}} \subseteq \mathcal{L}_{[0,1]}^\infty$, $q \in \mathcal{M}_\nu(\Theta, \Sigma)$, and T , Ψ_f^ν -regularized FTRL with regularizer scaling $\eta_t = \sqrt{c_2/t}$ achieves

$$R_T(q) \leq \sqrt{\frac{T+1}{c_2}} \Psi_f^\nu(dq/d\nu) + \frac{c_1}{\sqrt{c_2}} \sqrt{T}.$$

Proof of Corollary 1. We start by applying Theorem 2 using the intermediate point \tilde{z}_{t+1} and with $m_t \equiv 0$. Since $f'' \cdot (f + c_1) \geq c_2$ and either $1/f''$ or f is increasing, this simplifies to

$$\begin{aligned} R_T(q) & \leq \frac{1}{\eta_{T+1}} \Psi_f^\nu(dq/d\nu) + \frac{1}{2c_2} \sum_{t=1}^T \eta_t \int f(x_t(\theta)) (\ell_t(\theta))^2 \nu(d\theta) + \frac{c_1}{2c_2} \sum_{t=1}^T \eta_t \int (\ell_t(\theta))^2 \nu(d\theta) \\ & \quad - \sum_{t=1}^T \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \Psi_f^\nu(x_{t+1}). \end{aligned}$$

Taking $\eta_t = c/\sqrt{t}$ for some $c > 0$ gives

$$\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \geq \frac{1}{2c\sqrt{t+1}}.$$

Finally,

$$\sum_{t=1}^T \frac{1}{\sqrt{t}} \leq 1 + \int_1^T \frac{1}{\sqrt{x}} dx \leq 2\sqrt{T}.$$

Thus, using that $\ell_t \leq 1$ and $\Psi_f'(x_t) \geq 0$ (by Jensen's and $f(1) = 0$),

$$\begin{aligned} R_T(q) &\leq \frac{\sqrt{T+1}}{c} \Psi_f'(dq/d\nu) + \sum_{t=1}^T \frac{c}{2c_2\sqrt{t}} \Psi_f'(x_t) + \frac{c_1c}{c_2} \sqrt{T} - \sum_{t=1}^T \frac{1}{2c\sqrt{t+1}} \Psi_f'(x_{t+1}) \\ &\leq \frac{\sqrt{T+1}}{c} \Psi_f'(dq/d\nu) + \frac{c_1c}{c_2} \sqrt{T} + \sum_{t=2}^T \left(\frac{c}{2c_2\sqrt{t}} \Psi_f'(x_t) - \frac{1}{2c\sqrt{t}} \Psi_f'(x_t) \right), \end{aligned}$$

where the last step uses that $x_1 = \nu$. Taking $c = \sqrt{c_2}$ makes the final summation over t non-positive while minimizing the first term. \square

We also have the following result, which implies [Corollary 2](#).

Lemma 3. *If $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ and there exists $k_1 \in \mathbb{R}$ and $k_2 \in \mathbb{R}_+$ with $f(r) \leq k_1 + k_2 r \sqrt{\log(1+r)}$ for all $r \in \mathbb{R}_+$, then for all $\mu \in \mathcal{M}_\nu(\Theta, \Sigma)$*

$$D_f(\mu \| \nu) \leq k_1 + k_2 \sqrt{1 + \text{KL}(\mu \| \nu)}.$$

Proof of Lemma 3.

$$\begin{aligned} D_f(\mu \| \nu) &= \int f\left(\frac{d\mu}{d\nu}(\theta)\right) \nu(d\theta) \\ &\leq k_1 + k_2 \int \frac{d\mu}{d\nu}(\theta) \sqrt{1 + \log\left(\frac{d\mu}{d\nu}(\theta)\right)} \nu(d\theta) \\ &= k_1 + k_2 \int \sqrt{\log\left(1 + \frac{d\mu}{d\nu}(\theta)\right)} \mu(d\theta) \\ &\leq k_1 + k_2 \sqrt{\int \log\left(1 + \frac{d\mu}{d\nu}(\theta)\right) \mu(d\theta)} \\ &= k_1 + k_2 \sqrt{\int \frac{d\mu}{d\nu}(\theta) \log\left(1 + \frac{d\mu}{d\nu}(\theta)\right) \nu(d\theta)} \\ &\leq k_1 + k_2 \sqrt{1 + \text{KL}(\mu \| \nu)}. \end{aligned}$$

The first inequality is by assumption, the second is Jensen's, and the third is because $1 + (r \log r) \geq r \log(r+1)$ for all $r > 0$. \square

Finally, we restate and prove [Corollary 3](#).

Corollary 3. *Suppose $\nu \in \mathcal{M}^\infty(\Theta, \Sigma)$ and $f : \mathcal{S}(\nu) \rightarrow \mathbb{R}$ satisfies:*

1. $f(1/\nu(\Theta)) \geq 0$;
2. f is twice continuously differentiable, $f'' > 0$ on $\text{interior}(\mathcal{S}(\nu))$, and $f''(0+) > 0$.

For any sequence $(\ell_t)_{t \in \mathbb{N}} \subseteq \mathcal{L}_{[0,1]}^\infty$, $q \in \mathcal{M}_\nu(\Theta, \Sigma)$, and T , Ψ_f' -regularized FTRL achieves

1. *If $1/f''(x) \leq Cx$, $\eta_{t+1} = (C[1/2 + \sum_{s=1}^{t-1} \text{Var}_{\theta \sim \nu^{(z_{s+1})}} \ell_s(\theta)])^{-1/2}$ gives*

$$R_T(q) \leq [\Psi_f'(dq/d\nu) + 1] \sqrt{C[1/2 + \sum_{t=1}^T \text{Var}_{\theta \sim \nu^{(z_{t+1})}} \ell_t(\theta)]}.$$

2. *If $1/f'' \leq C$, $\eta_{t+1} = (C\nu(\Theta)[1/4 + \sum_{s=1}^t \text{Var}_{\theta \sim \bar{\nu}} \ell_s(\theta)])^{-1/2}$ for $\bar{\nu} = \nu/\nu(\Theta)$ gives*

$$R_T(q) \leq [\Psi_f'(dq/d\nu) + 1] \sqrt{C\nu(\Theta)[1/4 + \sum_{t=1}^T \text{Var}_{\theta \sim \bar{\nu}} \ell_t(\theta)]}.$$

Proof of Corollary 3. Note that, since $f(1/\nu(\Theta)) \geq 0$, Jensen's inequality implies that for any $x \in \mathcal{X}$

$$\Psi_f^\nu(x) = \int f(x(\theta))\nu(d\theta) \geq \nu(\Theta)f\left(\frac{1}{\nu(\Theta)} \int x(\theta)\nu(d\theta)\right) \geq 0.$$

First, suppose $1/f''(x) \leq Cx$. Since η_t is decreasing, [Theorem 2](#) with intermediate point A implies

$$R_T(q) \leq \frac{1}{\eta_{T+1}} \Psi_f^\nu\left(\frac{dq}{d\nu}\right) + \sum_{t=1}^T \frac{C\eta_t}{2} \int \hat{z}_{t+1}(\theta) (\ell_t(\theta) - m_t)^2 \nu(d\theta).$$

Taking $m_t = \int \ell_t(\theta)\hat{z}_{t+1}(\theta)\nu(d\theta)$ gives

$$\begin{aligned} R_T(q) &\leq \frac{1}{\eta_{T+1}} \Psi_f^\nu\left(\frac{dq}{d\nu}\right) + \frac{\sqrt{C}}{2} \sum_{t=1}^T \frac{\text{Var}_{\theta \sim \nu^{(\hat{z}_{t+1})}} \ell_t(\theta)}{\sqrt{1/2 + \sum_{s=1}^{t-2} \text{Var}_{\theta \sim \nu^{(\hat{z}_{s+1})}} \ell_s(\theta)}} \\ &\leq \frac{1}{\eta_{T+1}} \Psi_f^\nu\left(\frac{dq}{d\nu}\right) + \frac{\sqrt{C}}{2} \sum_{t=1}^T \frac{\text{Var}_{\theta \sim \nu^{(\hat{z}_{t+1})}} \ell_t(\theta)}{\sqrt{\sum_{s=1}^t \text{Var}_{\theta \sim \nu^{(\hat{z}_{s+1})}} \ell_s(\theta)}}, \end{aligned}$$

where we have used that the variance of a random variable in $[0, 1]$ is bounded by $1/4$ (see, e.g., Lemma 8 of [\[BNR20\]](#)). By Lemma 4.13 of [Orabona \[O19\]](#), this gives

$$R_T(q) \leq \Psi_f^\nu\left(\frac{dq}{d\nu}\right) \sqrt{C \left[1/2 + \sum_{t=1}^{T-1} \text{Var}_{\theta \sim \nu^{(\hat{z}_{t+1})}} \ell_t(\theta)\right]} + \sqrt{C \sum_{t=1}^T \text{Var}_{\theta \sim \nu^{(\hat{z}_{t+1})}} \ell_t(\theta)}.$$

Next, suppose $1/f''(\theta) \leq C$. Since η_t is decreasing, [Theorem 2](#) with either intermediate point implies

$$R_T(q) \leq \frac{1}{\eta_{T+1}} \Psi_f^\nu\left(\frac{dq}{d\nu}\right) + \sum_{t=1}^T \frac{\nu(\Theta)C\eta_t}{2} \int (\ell_t(\theta) - m_t)^2 \bar{\nu}(d\theta).$$

Taking $m_t = \int \ell_t(\theta)\bar{\nu}(\theta)$ gives

$$\begin{aligned} R_T(q) &\leq \frac{1}{\eta_{T+1}} \Psi_f^\nu\left(\frac{dq}{d\nu}\right) + \frac{\sqrt{\nu(\Theta)C}}{2} \sum_{t=1}^T \frac{\text{Var}_{\theta \sim \bar{\nu}} \ell_t(\theta)}{\sqrt{1/4 + \sum_{s=1}^{t-1} \text{Var}_{\theta \sim \bar{\nu}} \ell_s(\theta)}} \\ &\leq \frac{1}{\eta_{T+1}} \Psi_f^\nu\left(\frac{dq}{d\nu}\right) + \frac{\sqrt{\nu(\Theta)C}}{2} \sum_{t=1}^T \frac{\text{Var}_{\theta \sim \bar{\nu}} \ell_t(\theta)}{\sqrt{\sum_{s=1}^t \text{Var}_{\theta \sim \bar{\nu}} \ell_s(\theta)}} \end{aligned}$$

By Lemma 4.13 of [Orabona \[O19\]](#), this gives

$$R_T(q) \leq \Psi_f^\nu\left(\frac{dq}{d\nu}\right) \sqrt{\nu(\Theta)C \left[1/4 + \sum_{t=1}^T \text{Var}_{\theta \sim \bar{\nu}} \ell_t(\theta)\right]} + \sqrt{\nu(\Theta)C \sum_{t=1}^T \text{Var}_{\theta \sim \bar{\nu}} \ell_t(\theta)}.$$

□

C Proof of lower bound for quantile regret

Theorem 1. For all $N \in \mathbb{N}$ there exists a probability distribution p on $[0, 1]^N$ such that for any sequence of player predictions $(w_t)_{t \in \mathbb{N}} \subseteq \text{simp}([N])$, $i_\epsilon \in \{1, \dots, \lfloor N/4 \rfloor\}$, and $T \in \mathbb{N}$,

$$\mathbb{E}_{\ell_{1:T} \sim p^{\otimes T}} R_T(\delta_{(i_\epsilon)}) \geq \sqrt{(T/2)(\log(1/\epsilon) - 2 \log 2 + 1/\pi)} - \sqrt{2/\pi} - 2 \log N - \log 2,$$

where $\delta_{(i_\epsilon)}$ is the point-mass on the i_ϵ th best expert with respect to L_T and $\epsilon = i_\epsilon/N$.

Proof of Theorem 1. Let $\Phi(z) \stackrel{def}{=} \mathbb{P}(Z \leq z)$ for $Z \sim \text{Nor}(0, 1)$ be the normal cumulative distribution function and $\bar{\Phi}(z) \stackrel{def}{=} 1 - \Phi(z)$ be the normal complementary cumulative distribution function. Let $\ell_t(i) \sim \text{Ber}(1/2)$ for all $t \in [T]$ and $i \in [N]$. For a sequence of real values, $S = (s_j)_{j \in [m]}$, denote its empirical cumulative distribution function by $\hat{F}_S(x) \stackrel{def}{=} \sum_{j \in [m]} \mathbb{I}[x \leq s_j]$. For a non-decreasing càdlàg function F (right continuous and left limits exist), let its improper inverse be $F^+(y) = \inf \{x : F(x) \geq y\}$. Let $\tilde{L}_t = L_t - (t/2)\mathbb{1}_N$ be the centred cumulative losses. Clearly,

$$\mathbb{E}_{\ell_{1:T} \sim p^{\otimes T}} R_T(\delta_{(i_\epsilon)}) = -\mathbb{E} \hat{F}_{\tilde{L}_T}^+(\epsilon).$$

Without loss of generality, we can enrich our probability space so that there exists i.i.d. standard normal random variables $Z_T(i) \stackrel{iid}{\sim} \text{Nor}(0, 1)$ with $L_T(i) = F_{\text{Bin}(T, 1/2)}^+ \circ \Phi(Z_T(i))$. From Bretagnolle and Massart [BM89, Lemma 4] this coupling satisfies

$$\left| L_T(i) - T/2 - \frac{\sqrt{T}}{2} Z_T(i) \right| \leq 1 + \frac{Z_T(i)^2}{8},$$

which implies

$$\mathbb{E} \max_{i \in [N]} \left| L_T(i) - T/2 - \frac{\sqrt{T}}{2} Z_T(i) \right| \leq 1 + \frac{1}{8} \mathbb{E} \max_{i \in [N]} Z_T(i)^2.$$

Now,

$$\begin{aligned} \mathbb{E} \max_{i \in [N]} Z_T(i)^2 &\leq \inf_{\lambda \in (0, 1/2)} \frac{1}{\lambda} \log \sum_{i \in [N]} \mathbb{E} \exp(\lambda Z_T(i)^2) \\ &= \inf_{\lambda \in (0, 1/2)} \frac{1}{\lambda} \log \frac{N}{\sqrt{1-2\lambda}} \\ &= \inf_{\lambda \in (0, 1/2)} \left(\frac{\log N}{\lambda} + \frac{1}{\lambda} \log \frac{1}{\sqrt{1-2\lambda}} \right). \end{aligned}$$

Overapproximating with $\lambda = 1/4$ gives

$$\mathbb{E} \max_{i \in [N]} Z_T(i)^2 \leq 4 \log N + 2 \log 2.$$

From Ali and Chan [AC65], for $\epsilon > 3/4$

$$\begin{aligned} -\mathbb{E} \hat{F}_{Z_T}^+(\epsilon) &= \mathbb{E} \hat{F}_{-Z_T}^+(1-\epsilon) \\ &\geq \mathbb{E} \Phi^+ \left(\frac{\lfloor (1-\epsilon)N \rfloor}{N+1} \right) \\ &= \mathbb{E} \bar{\Phi}^{-1} \left(1 - \frac{\lfloor (1-\epsilon)N \rfloor}{N+1} \right). \end{aligned}$$

Then, since the coupling is monotone, transformations of the order statistics are equal to the order statistics of the transformation, and hence

$$\begin{aligned} &\left| \hat{F}_{-Z_T}^+(1-\epsilon) - \hat{F}_{\left(-\frac{2}{\sqrt{T}}(L_T-T/2)\right)}^+(1-\epsilon) \right| \\ &\leq \max_{i \in [N]} \left| Z_T(i) - \left(\frac{2}{\sqrt{T}} (L_T(i) - T/2) \right) \right|. \end{aligned}$$

Thus

$$\mathbb{E} \left| \hat{F}_{-Z_T}^+(1-\epsilon) - \hat{F}_{\left(-\frac{2}{\sqrt{T}}(L_T-T/2)\right)}^+(1-\epsilon) \right| \leq \frac{4 \log N + 2 \log 2}{\sqrt{T}},$$

so

$$\mathbb{E} \hat{F}_{\left(-\frac{2}{\sqrt{T}}(L_T-T/2)\right)}^+(1-\epsilon) \geq \bar{\Phi}^{-1} \left(1 - \frac{(1-\epsilon)N}{N+1} \right) - \frac{4 \log N + 2 \log 2}{\sqrt{T}},$$

which implies

$$\begin{aligned} -\mathbb{E}\widehat{F}_{L_T}^+(\epsilon) &\geq \frac{\sqrt{T}}{2}\overline{\Phi}^{-1}\left(\frac{N}{N+1}\epsilon + \frac{1}{N+1}\right) - 2\log N - \log 2 \\ &\geq \frac{\sqrt{T}}{2}\overline{\Phi}^{-1}(2\epsilon) - 2\log N - \log 2. \end{aligned}$$

Now, applying [Lemma 4](#), for $\epsilon < 1/4$

$$-\mathbb{E}\widehat{F}_{L_T}^+(\epsilon) \geq \sqrt{\frac{T}{2}(\log(1/\epsilon) - 2\log 2 + 1/\pi)} - \sqrt{2/\pi} - 2\log N - \log 2.$$

□

Lemma 4. For $x > 0$,

$$\overline{\Phi}(x) \geq \frac{\exp(1/\pi)\exp(-(x + \sqrt{2/\pi})^2/2)}{2},$$

and for $0 < y < 1/2$,

$$\overline{\Phi}^{-1}(y) \geq \sqrt{2\log(1/y) - 2\log 2 + 2/\pi} - \sqrt{2/\pi}.$$

Proof. The second bound follows from the first, so we will only prove the first. Let $\phi(x)$ denote $\overline{\Phi}'(x) = \exp(-x^2/2)/\sqrt{2\pi}$, and

$$\begin{aligned} h(x) &= \log \overline{\Phi}(x) - \log \frac{\exp(1/\pi)\exp(-(x + \sqrt{2/\pi})^2/2)}{2} \\ &= \log \overline{\Phi}(x) - 1/\pi + \log(2) + (x + \sqrt{2/\pi})^2/2. \end{aligned}$$

Then, $h(0) = 0$ and

$$h'(x) = -\phi(x)/\overline{\Phi}(x) + (x + \sqrt{2/\pi}),$$

so $h'(0) = 0$ as well.

Next,

$$h''(x) = x\phi(x)/\overline{\Phi}(x) - (\phi(x)/\overline{\Phi}(x))^2 + 1,$$

so $h''(0) = 1 - 2/\pi > 0$.

Suppose towards a contradiction that $h''(x_0) = 0$ for some $x_0 > 0$. Rearranging gives

$$x_0 = \phi(x_0)/\overline{\Phi}(x_0) - (\phi(x_0)/\overline{\Phi}(x_0))^{-1},$$

so

$$\phi(x_0)/\overline{\Phi}(x_0) = \frac{\sqrt{4 + x_0^2} + x_0}{2}.$$

However, from [Birnbaum \[B42\]²](#), for all $x > 0$

$$\phi(x)/\overline{\Phi}(x) < \frac{\sqrt{4 + x^2} + x}{2}.$$

Thus $h''(x) \neq 0$ for all $x > 0$. Since h'' is continuous, $h''(0) > 0$, and h'' has no zeros on $(0, \infty)$, by the intermediate value theorem $h''(x) > 0$ for all $x > 0$. Finally, since $h(0) = h'(0) = 0$ and $h''(x) > 0$ for all $x > 0$, by standard differential inequalities, $h(x) \geq 0$ for all $x \geq 0$. □

²The stated result has a non-strict inequality, however the proof by Jensen's inequality clearly leads to a strict inequality since the function in the application of Jensen's inequality, $t \mapsto 1/t$ is strictly convex and the distribution to which it is applied is not degenerate, it has density proportional to $\phi'(t)$ for $t > 0$.

D Proof of semi-adversarial regret bound

To analyze FTRL-CARL, we define two more notions to characterize which stage of learning the algorithm is in. Note that each of these stages occur *implicitly* while running FTRL-CARL, and consequently the user does not need to know any of the quantities associated with \mathcal{D} .

For each $i \in [N] \setminus \mathcal{I}_0$, let $T_i = \lceil 8(\log N)/\Delta_i^2 \rceil$ and $T_0 = \max_{i \in [N] \setminus \mathcal{I}_0} T_i$. For each $t \in \mathbb{N}$ let $\mathcal{I}_0^{(t)} = \mathcal{I}_0 \cup \{i \in [N] \setminus \mathcal{I}_0 \text{ s.t. } T_i \geq t\}$ and $N_0^{(t)} = |\mathcal{I}_0^{(t)}|$. For each $j \in \{0, 1, \dots, N - N_0 - 1\}$, let $T_{(j)}$ be the decreasing ordered values of T_i , so that $T_0 = T_{(0)} \geq T_{(1)} \geq \dots \geq T_{(N-N_0-1)}$. Let $T_{(N-N_0)} = 0$. Let $\Delta_{(j)}$ be the corresponding increasing ordered values of Δ_i , so that $\Delta_0 = \Delta_{(0)} \leq \Delta_{(1)} \leq \dots \leq \Delta_{(N-N_0-1)}$.

We then have the following regret bound for FTRL-CARL, which simplifies to [Corollary 4](#) by lower bounding all Δ_i by Δ_0 and observing that W_{j,N,N_0} telescopes in the sum. The more refined bound of [Theorem 3](#) will be significantly tighter for small T when only one Δ_i is small, although the worst-case adversary will choose all $\Delta_i = \Delta_0$.

Theorem 3. *For any time-homogeneous convex constraint \mathcal{D} , FTRL-CARL achieves:*

For all T ,

$$\mathbb{E}R_T \leq \sqrt{2T \log N},$$

and if $T > T_0$,

$$\begin{aligned} \mathbb{E}R_T \leq & \sqrt{2T \log N_0} + 4(\log N) \sum_{j=0}^{N-N_0-1} W_{j,N,N_0} \frac{1}{\Delta_{(j)}} \\ & + \frac{5\sqrt{2}}{N\sqrt{\log N}} (e^{-1/2} + \mathbb{I}[N_0 = 1]) \sum_{i \in [N] \setminus \mathcal{I}_0} \frac{\mathbb{I}[T > T_i]}{\Delta_i} + \sqrt{\log N}, \end{aligned}$$

$$\text{where } W_{j,N,N_0} = \frac{1}{\sqrt{\log N}} \left(\sqrt{\log(N_0 + j + 1)} - \sqrt{\log(N_0 + j)} \right).$$

Proof of Theorem 3. Since $N < \infty$, we wish to apply [Lemma 1](#) and [Theorem 2](#) for ν equal to counting measure and $\mathcal{S}(\nu) = [0, 1]$. It is straightforward to verify that $-h''_C(x) = 1/h_A(x) > 0$ on $(0, 1)$. Thus, we can apply [Theorem 2](#) with intermediate point A , so there exists a sequence $w_{t+1} \in \text{Conv}(w_t^{\text{CARL}}, w_{t+1}^{\text{CARL}})$ such that for any sequence $(m_t)_{t \in \mathbb{N}} \subseteq \mathbb{R}$, it holds almost surely that

$$R_T \leq \sqrt{\log N} + \sum_{t=1}^T \sum_{i \in [N]} \left[\frac{c}{2\sqrt{t}} h_A(v_{t+1}(i)) (\ell_t(i) - m_t)^2 + \frac{1}{c} (\sqrt{t+1} - \sqrt{t}) H_C(w_{t+1}^{\text{CARL}}(i)) \right],$$

where we have used that $H_C(w) = 0$ for any one-hot w , an application of [Lemma 9](#), and the expression for h'' again.

Let $h_A(x) = x\sqrt{2\log(1/x)}$ for $x \in (0, 1]$ and $h_A(0) = 0$, and set $H_A(w) = \sum_{i \in [N]} h_A(w(i))$. Define $\gamma : \text{simp}([N]) \rightarrow \text{simp}([N])$ by

$$\gamma(w) = \begin{cases} (1/N, \dots, 1/N), & H_A(w) = 0 \\ \left(\frac{h_A(w(i))}{H_A(w)} \right)_{i \in [N]}, & \text{otherwise.} \end{cases}$$

Letting $m_t = \langle \ell_t, \gamma(v_{t+1}) \rangle$ then gives

$$R_T \leq \sqrt{\log N} + \sum_{t=1}^T \left[\frac{c}{2\sqrt{t}} H_A(v_{t+1}) \text{Var}_{I \sim \gamma(v_{t+1})} [\ell_t(I)] + \frac{1}{c} (\sqrt{t+1} - \sqrt{t}) H_C(w_{t+1}^{\text{CARL}}) \right],$$

Since $0 \leq \ell_t \leq 1$, $\text{Var}_{I \sim \gamma(w)} [\ell_t(I)] \leq 1/4$ for any $w \in \text{simp}([N])$, which combined with $\sqrt{t+1} - \sqrt{t} \leq 1/(2\sqrt{t})$ gives

$$R_T \leq \sqrt{\log N} + \sum_{t=1}^T \left[\frac{c}{8\sqrt{t}} H_A(v_{t+1}) + \frac{1}{2c\sqrt{t}} H_C(w_{t+1}^{\text{CARL}}) \right]. \quad (14)$$

Applying Lemma 9 to Eq. (14) with any p and $\mathcal{I}_0 = [N]$ gives the worst-case bound of

$$R_T \leq \sqrt{\log N} + \left(\frac{c}{4} + \frac{1}{c}\right) \sum_{t=1}^T \frac{1}{2\sqrt{t}} \sqrt{2 \log N} \leq \left(\frac{c}{4} + \frac{1}{c}\right) \sqrt{2T \log N}.$$

Next, applying Lemma 9 to Eq. (14) with $p = 1/2$ and $\mathcal{I}_0 = \mathcal{I}_0^{(t)}$,

$$\begin{aligned} R_T &\leq \sqrt{\log N} + \left(\frac{c}{4} + \frac{1}{c}\right) \sum_{t=1}^T \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} \\ &\quad + \sum_{t=1}^T \frac{c}{8\sqrt{t}} \left[\frac{1}{\sqrt{e/2}} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{v_{t+1}(i)} + \mathbb{I}[N_0^{(t)} = 1] \sqrt{2} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{v_{t+1}(i)} \right] \\ &\quad + \sum_{t=1}^T \frac{1}{2c\sqrt{t}} \left[\frac{1}{\sqrt{e/2}} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{w_{t+1}^{\text{CARL}}(i)} + \mathbb{I}[N_0^{(t)} = 1] \sqrt{2} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{w_{t+1}^{\text{CARL}}(i)} \right]. \end{aligned}$$

Since $v_{t+1} \in \text{Conv}(w_t^{\text{CARL}}, w_{t+1}^{\text{CARL}})$, then $\sqrt{v_{t+1}} \leq \sqrt{w_t^{\text{CARL}}} + \sqrt{w_{t+1}^{\text{CARL}}}$, and so

$$\begin{aligned} R_T &\leq \sqrt{\log N} + \left(\frac{c}{4} + \frac{1}{c}\right) \sum_{t=1}^T \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} \\ &\quad + \sum_{t=1}^T \frac{c}{8\sqrt{t}} \left[\frac{1}{\sqrt{e/2}} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{w_t^{\text{CARL}}(i)} + \mathbb{I}[N_0^{(t)} = 1] \sqrt{2} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{w_t^{\text{CARL}}(i)} \right] \\ &\quad + \sum_{t=1}^T \left(\frac{c}{8\sqrt{t}} + \frac{1}{2c\sqrt{t}}\right) \left[\frac{1}{\sqrt{e/2}} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{w_{t+1}^{\text{CARL}}(i)} + \mathbb{I}[N_0^{(t)} = 1] \sqrt{2} \sum_{i \in [N] \setminus \mathcal{I}_0^{(t)}} \sqrt{w_{t+1}^{\text{CARL}}(i)} \right]. \end{aligned}$$

Simplifying gives

$$\begin{aligned} R_T &\leq \sqrt{\log N} + \left(\frac{c}{4} + \frac{1}{c}\right) \sum_{t=1}^T \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} \\ &\quad + \frac{1}{\sqrt{e/2}} \sum_{i \in [N] \setminus \mathcal{I}_0} \sum_{t=T_{i+1}}^T \frac{1}{2\sqrt{t}} \left[\frac{c}{4} \sqrt{w_t^{\text{CARL}}(i)} + \left(\frac{1}{c} + \frac{c}{4}\right) \sqrt{w_{t+1}^{\text{CARL}}(i)} \right] \\ &\quad + \mathbb{I}[N_0 = 1] \sqrt{2} \sum_{i \in [N] \setminus \mathcal{I}_0} \sum_{t=T_0+1}^T \frac{1}{2\sqrt{t}} \left[\frac{c}{4} \sqrt{w_t^{\text{CARL}}(i)} + \left(\frac{1}{c} + \frac{c}{4}\right) \sqrt{w_{t+1}^{\text{CARL}}(i)} \right]. \end{aligned}$$

For $T > T_0$, the first term can be decomposed as

$$\begin{aligned} \sum_{t=1}^T \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} &= \sum_{t=1}^{T_0} \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} + \sum_{t=T_0+1}^T \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0} \\ &\leq \sum_{t=1}^{T_0} \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} + \sqrt{2T \log N_0} - \sqrt{2T_0 \log N_0}. \end{aligned}$$

For the second and third terms, applying [Lemma 8](#), for each $i \in [N] \setminus \mathcal{I}_0$ we have

$$\begin{aligned} & \mathbb{E} \sum_{t=T_i+1}^T \frac{1}{2\sqrt{t}} \left[\frac{c}{4} \sqrt{w_t^{\text{CARL}}(i)} + \left(\frac{1}{c} + \frac{c}{4} \right) \sqrt{w_{t+1}^{\text{CARL}}(i)} \right] \\ & \leq \frac{c}{8} \sum_{t=T_i+1}^T \frac{1}{\sqrt{t}} \mathbb{E} [(w_t^{\text{CARL}}(i))^p] + \left(\frac{1}{2c} + \frac{c}{8} \right) \sum_{t=T_i+2}^T \frac{2}{\sqrt{t}} \mathbb{E} [(w_t^{\text{CARL}}(i))^p] \\ & \leq \left(\frac{1}{c} + \frac{3c}{8} \right) \frac{\sqrt{8}(2+c)}{N\Delta_i c \sqrt{\log N}}. \end{aligned}$$

Taking $c = 2$ and combining these results, we have

$$\begin{aligned} \mathbb{E} R_T & \leq \left(\sqrt{2T \log N_0} - \sqrt{2T_0 \log N_0} \right) \mathbb{I}[T > T_0] + \sum_{t=1}^{\min(T, T_0)} \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} \\ & \quad + \frac{5\sqrt{2}}{N\sqrt{\log N}} \left(e^{-1/2} + \mathbb{I}[N_0 = 1] \right) \sum_{i \in [N] \setminus \mathcal{I}_0} \frac{\mathbb{I}[T > T_i]}{\Delta_i} + \sqrt{\log N}. \end{aligned}$$

Next, using summation by parts,

$$\begin{aligned} & \sum_{t=1}^{T_0} \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} \\ & = \sum_{j=0}^{N-N_0} \sum_{t=T_{(N-N_0-j)+1}}^{T_{(N-N_0-j+1)}} \frac{1}{2\sqrt{t}} \sqrt{2 \log(N-j)} \\ & \leq \sum_{j=0}^{N-N_0} \left(\sqrt{T_{(N-N_0-j+1)}} - \sqrt{T_{(N-N_0-j)+1}} \right) \sqrt{2 \log(N-j)} \\ & = \sqrt{2T_{(1)} \log N_0} - \sqrt{2T_{(N-N_0)} \log N} \\ & \quad + \sum_{j=1}^{N-N_0} \sqrt{T_{(N-N_0-j)}} \left(\sqrt{2 \log(N-j+1)} - \sqrt{2 \log(N-j)} \right) \\ & \leq \sqrt{2T_0 \log N_0} + \sum_{j=0}^{N-N_0-1} \sqrt{T_{(j)}} \left(\sqrt{2 \log(N_0+j+1)} - \sqrt{2 \log(N_0+j)} \right) \\ & \leq \sqrt{2T_0 \log N_0} + 4\sqrt{\log N} \sum_{j=0}^{N-N_0-1} \frac{1}{\Delta_{(j)}} \left(\sqrt{\log(N_0+j+1)} - \sqrt{\log(N_0+j)} \right). \end{aligned}$$

In summary, for all $T > 0$

$$\begin{aligned} \mathbb{E} R_T & \leq \sum_{t=1}^T \frac{1}{2\sqrt{t}} \sqrt{2 \log N_0^{(t)}} + \sqrt{\log N} \\ & \quad + \frac{5\sqrt{2}}{N\sqrt{\log N}} \left(e^{-1/2} + \mathbb{I}[N_0 = 1] \right) \sum_{i \in [N] \setminus \mathcal{I}_0} \frac{\mathbb{I}[T > T_i]}{\Delta_i}, \end{aligned}$$

and for $T > T_0$,

$$\begin{aligned} \mathbb{E} R_T & \leq \sqrt{2T \log N_0} + 4(\log N) \sum_{j=0}^{N-N_0-1} W_{j, N, N_0} \frac{1}{\Delta_{(j)}} \\ & \quad + \frac{5\sqrt{2}}{N\sqrt{\log N}} \left(e^{-1/2} + \mathbb{I}[N_0 = 1] \right) \sum_{i \in [N] \setminus \mathcal{I}_0} \frac{\mathbb{I}[T > T_i]}{\Delta_i} + \sqrt{\log N}, \end{aligned}$$

where

$$W_{j,N,N_0} = \frac{\sqrt{\log(N_0 + j + 1)} - \sqrt{\log(N_0 + j)}}{\sqrt{\log N}}.$$

Note that $\sum_{j=0}^{N-N_0-1} W_{j,N,N_0} \leq 1$. \square

D.1 Bounding the weights of ineffective experts

Let $I_t^* = \arg \min_{i \in [N]} L_t(i)$, so $\delta_{(1)}$ is the one-hot vector with a one on I_t^* . Then, we have the following control on the weights.

Lemma 5. For any $i \in [N]$, w_{t+1}^{CARL} satisfies

$$w_{t+1}^{\text{CARL}}(i) \leq \exp \left\{ -c^2 \frac{(L_t(i) - L_t(I_t^*))^2}{2(t+1)} \right\}.$$

Proof of Lemma 5. The Lagrangian corresponding to the optimization problem defining w_{t+1}^{CARL} is

$$\mathcal{L}(w; \lambda, \alpha_1, \dots, \alpha_N) = \langle L_t, w \rangle - \frac{\sqrt{t+1}}{c} H_C(w) + \lambda(1 - \langle \mathbf{1}, w \rangle) + \sum_{i \in [N]} \alpha_i w(i).$$

It is straightforward to verify that $h'_C(x) = \sqrt{2 \log(1/x)}$, so

$$\frac{\partial}{\partial w(i)} \mathcal{L}(w; \lambda, \alpha_1, \dots, \alpha_N) = L_t(i) - \frac{\sqrt{t+1}}{c} \sqrt{2 \log(1/w(i))} - \lambda + \alpha_i.$$

Setting this to 0 gives

$$\lambda^* = L_t(i) - \frac{\sqrt{t+1}}{c} \sqrt{2 \log(1/w(i))} + \alpha_i.$$

Since H_C is invariant to reordering of the expert indices, w_{t+1}^{CARL} is monotonically non-decreasing in L_t , so $w_{t+1}^{\text{CARL}}(I_t^*) \geq 1/N$ and consequently $\alpha_{I_t^*} = 0$. If $w_{t+1}^{\text{CARL}}(i) = 0$, the statement of the lemma trivially holds. Otherwise, $\alpha_i = 0$, so by definition of λ^* we have

$$\sqrt{2 \log(1/w_{t+1}^{\text{CARL}}(i))} - \sqrt{2 \log(1/w_{t+1}^{\text{CARL}}(I_t^*))} = \frac{c}{\sqrt{t+1}} (L_t(i) - L_t(I_t^*)).$$

The result follows by rearranging this equation and using that $\sqrt{a} - \sqrt{b} \leq \sqrt{a-b}$ for all $a \geq b \geq 0$ and $w_{t+1}^{\text{CARL}}(i) \leq w_{t+1}^{\text{CARL}}(I_t^*)$. \square

Next, we extend Theorem 1 of Bilodeau, Negrea, and Roy [BNR20] to depend on the individual ineffective experts effective gaps.

Lemma 6. For all $N \geq 2$, convex sets $\mathcal{D} \subseteq \mathcal{M}(\hat{\mathcal{Y}}^N \times \mathcal{Y})$, $\lambda > 0$, t , and $i \in [N] \setminus \mathcal{I}_0$, if the environment is subject to the time-homogeneous convex constraint then

$$\mathbb{E} \exp \left\{ \lambda (L_t(I_t^*) - L_t(i)) \right\} \leq \exp \left\{ T [\lambda^2/2 - \lambda \Delta_i] \right\}.$$

Proof of Lemma 6. The argument follows nearly identically to the one given by Bilodeau, Negrea, and Roy [BNR20]. However, in the second last step in their argument, note that

$$\inf_{w \in \text{simp}(\mathcal{I}_0)} \mathbb{E}_\ell \exp \left\{ \lambda [\langle \ell(\mathcal{I}_0), w \rangle - \ell(i)] \right\} \leq e^{\lambda^2/2 - \lambda \Delta_i}.$$

The last step is then applied in exactly the analogous way. \square

Using these two results, we control the expected size of the weight of an ineffective expert.

Lemma 7. For every t , $i \in [N] \setminus \mathcal{I}_0$, and $p > 0$

$$\mathbb{E} [(w_t^{\text{CARL}}(i))^p] \leq 2 \exp \left\{ -t \frac{pc^2}{2(\sqrt{2} + c\sqrt{p})^2} \Delta_i^2 \right\}.$$

Proof of Lemma 7. For any $s \in (0, 1)$,

$$\mathbb{E} [(w_t^{\text{CARL}}(i))^p] \leq \mathbb{E} [(w_t^{\text{CARL}}(i))^p : L_t(i) - L_t(I_t^*) \geq s\Delta_i t] + \mathbb{P}[L_t(i) - L_t(I_t^*) < s\Delta_i t].$$

From Lemma 5,

$$\begin{aligned} & \mathbb{E} [(w_t^{\text{CARL}}(i))^p : L_t(i) - L_t(I_t^*) \geq s\Delta_i t] \\ & \leq \mathbb{E} \left[\exp \left\{ -pc^2 \frac{(L_t(i) - L_t(I_t^*))^2}{2(t+1)} \right\} : L_t(i) - L_t(I_t^*) \geq s\Delta_i t \right] \\ & \leq \exp \left\{ -ps^2 c^2 \Delta_i^2 \frac{t^2}{2(t+1)} \right\} \\ & \leq \exp \left\{ -tps^2 c^2 \Delta_i^2 / 4 \right\}. \end{aligned}$$

Next, using the Cramèr-Chernoff Method and Lemma 6,

$$\begin{aligned} & \mathbb{P}[L_t(i) - L_t(I_t^*) < s\Delta_i t] \\ & = \mathbb{P}[L_t(I_t^*) - L_t(i) > -s\Delta_i t] \\ & \leq \inf_{\lambda > 0} \frac{\mathbb{E} \exp \{ \lambda (L_t(I_t^*) - L_t(i)) \}}{\exp \{ -\lambda s \Delta_i t \}} \\ & \leq \inf_{\lambda > 0} \frac{\exp \{ t (\lambda^2 / 2 - \lambda \Delta_i) \}}{\exp \{ -\lambda s \Delta_i t \}} \\ & = \inf_{\lambda > 0} \exp \{ t (\lambda^2 / 2 - (1-s)\lambda \Delta_i) \} \\ & = \exp \{ -t (1-s)^2 \Delta_i^2 / 2 \}, \end{aligned}$$

where the last step follows from taking $\lambda = \Delta_i(1-s)$. These two terms decay at the same rate if $ps^2 c^2 = 2(1-s)$; that is, $s = \sqrt{2}/(\sqrt{2} + c\sqrt{p})$. Using this value,

$$\mathbb{E} [(w_t^{\text{CARL}}(i))^p] \leq 2 \exp \left\{ -t \frac{pc^2}{2(\sqrt{2} + c\sqrt{p})^2} \Delta_i^2 \right\}.$$

□

This single-round bound can now be summed to compute the overall regret contribution of ineffective experts.

Lemma 8. For every $i \in [N] \setminus \mathcal{I}_0$ and $p > 0$,

$$\sum_{t=T_i+1}^{\infty} \frac{1}{\sqrt{t}} \mathbb{E} [(w_t^{\text{CARL}}(i))^p] \leq \frac{2\sqrt{2}(\sqrt{2} + c\sqrt{p})}{N\Delta_i c\sqrt{p} \log N}.$$

Proof of Lemma 8. Using Lemma 7, we have

$$\begin{aligned} \sum_{t=T_i+1}^{\infty} \frac{1}{\sqrt{t}} \mathbb{E} [(w_t^{\text{CARL}}(i))^p] & \leq 2 \sum_{t=T_i+1}^{\infty} \frac{1}{\sqrt{t}} \exp \left\{ -t \frac{pc^2}{2(\sqrt{2} + c\sqrt{p})^2} \Delta_i^2 \right\} \\ & \leq 2 \int_{T_i}^{\infty} \frac{1}{\sqrt{t}} \exp \left\{ -t \frac{pc^2}{2(\sqrt{2} + c\sqrt{p})^2} \Delta_i^2 \right\} dt. \end{aligned}$$

For any $C > 0$, substituting $r^2/2 = tC$ and using $\mathbb{P}(Z > a) \leq \frac{\exp(-a^2/2)}{a\sqrt{2\pi}}$ for $Z \sim \text{Nor}(0, 1)$ gives

$$\begin{aligned} \int_{T_i}^{\infty} \frac{1}{\sqrt{t}} \exp \{ -tC \} dt & = \sqrt{2/C} \int_{\sqrt{2CT_i}}^{\infty} \exp \{ -r^2/2 \} dr \\ & \leq \sqrt{2/C} \frac{\exp \{ -CT_i \}}{\sqrt{2CT_i}} \\ & = \frac{\exp \{ -CT_i \}}{C\sqrt{T_i}}. \end{aligned}$$

Thus,

$$\begin{aligned} \sum_{t=T_i+1}^{\infty} \frac{1}{\sqrt{t}} \mathbb{E} [(w_t^{\text{CARL}}(i))^p] &\leq 2 \frac{\exp \left\{ -T_i \frac{pc^2}{2(\sqrt{2}+c\sqrt{p})^2} \Delta_i^2 \right\}}{\frac{pc^2}{2(\sqrt{2}+c\sqrt{p})^2} \Delta_i^2 \sqrt{T_i}} \\ &\leq \frac{2\sqrt{2}(\sqrt{2}+c\sqrt{p})}{N \Delta_i c \sqrt{p} \log N}. \end{aligned}$$

□

D.2 Entropic concentration of FTRL-CARL regularizer

We now show that the functions used to define the regularizer for w^{CARL} concentrate like an entropy. This result is the analogue of Lemma 1 of Bilodeau, Negrea, and Roy [BNR20].

Lemma 9. *For every $w \in \text{simp}([N])$, $\mathcal{I}_0 \subseteq [N]$ with $|\mathcal{I}_0| = N_0$, and $p \in (0, 1)$,*

$$0 \leq H_C(w) \leq H_A(w) \leq \sqrt{2 \log N_0} + \frac{1}{\sqrt{e(1-p)}} \sum_{i \in [N] \setminus \mathcal{I}_0} w(i)^p + \mathbb{I}[N_0 = 1] \sqrt{2} \sum_{i \in [N] \setminus \mathcal{I}_0} \sqrt{w(i)}.$$

Proof of Lemma 9. Using the previously observed expressions for the first and second derivatives, h_C is concave on $(0, 1)$, so H_C is concave on $\text{simp}([N])$. Bauer's maximum principle [B60] says that a concave function on a convex, compact set S attains its minimum at an extreme point of S . This means that H_C attains its minimum in the set of extreme points of $\text{simp}([N])$, which corresponds to the set of standard basis vectors. In particular, letting e_i denote the i th basis vector,

$$H_C(w) \geq \min_{i \in [N]} H_C(e_i) = 0.$$

For the second inequality, let

$$g(x) = \begin{cases} \sqrt{\frac{\pi}{2}} \operatorname{erf} \left(\sqrt{\log(1/x)} \right) - x(N-1) \sqrt{\frac{\pi}{2}}, & x \in (0, 1] \\ \sqrt{\pi/2}, & x = 0. \end{cases}$$

so that

$$H_C(w) = H_A(w) - \sum_{i \in [N]} g(w(i)).$$

Then,

$$\frac{d^2}{dx^2} g(x) = -\frac{1}{2\sqrt{2}x \log^{3/2}(1/x)}.$$

This is negative for $x \in (0, 1)$, so $\sum_{i \in [N]} g(w(i))$ is concave. Again, by Bauer's maximum principle, we have

$$\sum_{i \in [N]} g(w(i)) \geq 0,$$

so

$$H_C(w) \leq H_A(w).$$

For the third inequality, observe that

$$H_A(w) = \sum_{i_0 \in \mathcal{I}_0} w(i_0) \sqrt{2 \log(1/w(i_0))} + \sum_{i \in [N] \setminus \mathcal{I}_0} w(i) \sqrt{2 \log(1/w(i))}.$$

Consider the convex optimization problem

$$\min_{\substack{v \in \mathbb{R}_+^{N_0} \\ \langle \mathbf{1}, v \rangle \leq 1}} \left\{ - \sum_{i_0 \in \mathcal{I}_0} v(i_0) \sqrt{2 \log(1/v(i_0))} \right\}.$$

This has the Lagrangian

$$\mathcal{L}(v; \alpha, \lambda_1, \dots, \lambda_{N_0}) = - \sum_{i_0 \in \mathcal{I}_0} v(i_0) \sqrt{2 \log(1/v(i_0))} + \alpha (\langle \mathbf{1}, v \rangle - 1) + \sum_{i_0 \in \mathcal{I}_0} \lambda_{i_0} v(i_0).$$

The partial derivatives for $i_0 \in \mathcal{I}_0$ are

$$\partial_{i_0} \mathcal{L}(v; \alpha, \lambda_1, \dots, \lambda_{N_0}) = -\sqrt{2 \log(1/w(i_0))} + \frac{1}{\sqrt{2 \log(1/w(i_0))}} + \alpha + \lambda_{i_0}.$$

Note that this is undefined at $w(i_0) = 0$, so $\lambda_{i_0} = 0$ for all i_0 . If $\alpha = 0$, then for each $i_0 \in \mathcal{I}_0$, $w(i_0) = e^{-1/2}$, and this is only feasible when $N_0 = 1$. In this case,

$$\sum_{i_0 \in \mathcal{I}_0} w(i_0) \sqrt{2 \log(1/w(i_0))} = \frac{1}{\sqrt{e}}.$$

Otherwise, $\alpha > 0$, and since $-\sqrt{2 \log(1/x)} + \frac{1}{\sqrt{2 \log(1/x)}}$ is monotonic in x , all the $w(i_0)$ are equal.

By the K.K.T. condition, this implies $w(i_0) = 1/N_0$ for each i_0 , so

$$\sum_{i_0 \in \mathcal{I}_0} w(i_0) \sqrt{2 \log(1/w(i_0))} = \sqrt{2 \log N_0}.$$

That is, if $N_0 \geq 2$,

$$H_A(w) \leq \sqrt{2 \log N_0} + \sum_{i \in [N] \setminus \mathcal{I}_0} w(i) \sqrt{2 \log(1/w(i))}. \quad (15)$$

Next, if $\mathcal{I}_0 = \{i_0\}$, then we have

$$\begin{aligned} H_A(w) &= w(i_0) \sqrt{2 \log(1/w(i_0))} + \sum_{i \in [N] \setminus \mathcal{I}_0} w(i) \sqrt{2 \log(1/w(i))} \\ &\leq \sqrt{2w(i_0) \log(1/w(i_0))} + \sum_{i \in [N] \setminus \mathcal{I}_0} w(i) \sqrt{2 \log(1/w(i))} \\ &\leq \sqrt{2(1-w(i_0))} + \sum_{i \in [N] \setminus \mathcal{I}_0} w(i) \sqrt{2 \log(1/w(i))} \\ &= \sqrt{2 \sum_{i \in [N] \setminus \mathcal{I}_0} w(i)} + \sum_{i \in [N] \setminus \mathcal{I}_0} w(i) \sqrt{2 \log(1/w(i))} \\ &\leq \sum_{i \in [N] \setminus \mathcal{I}_0} \left[\sqrt{2w(i)} + w(i) \sqrt{2 \log(1/w(i))} \right], \end{aligned} \quad (16)$$

where we have used $x \log(1/x) \leq 1 - x$ for all $x > 0$ and that $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ for all $x, y > 0$.

Finally, for any $p \in (0, 1)$, let $f(x) = x^{1-p} \sqrt{2 \log(1/x)}$. Observe that $f(0^+) = 0$, $f(1) = 0$, and

$$f'(x) = (1-p)x^{-p} \sqrt{2 \log(1/x)} - \frac{x^{-p}}{\sqrt{2 \log(1/x)}}.$$

Thus, the only critical point of f occurs at $x_0 = e^{-\frac{1}{2(1-p)}}$. Substituting this gives

$$f(x) \leq \frac{1}{\sqrt{e(1-p)}}.$$

It follows that for $x \in [0, 1]$

$$x \sqrt{2 \log(1/x)} \leq \frac{x^p}{\sqrt{e(1-p)}} \quad (17)$$

Combining Eqs. (15) to (17) gives the third inequality. \square

E Experiments

In this section, we provide an empirical evaluation of our proposed algorithms and compare them to existing algorithms in the respective paradigms. In particular, for the quantile regret paradigm we implement the version of abNormal proposed in [Example 3](#), while for the semi-adversarial paradigm we implement FTRL-CARL and compare it to Meta-CARE [\[BNR20\]](#) and Hedge with a decreasing regularizer scaling [\[MG19\]](#). In both cases we generate synthetic data, with specific details in the respective sections. All algorithms use standard tuning for hyperparameters; the specific details can be found in the code, available at <https://github.com/blairbilodeau/neurips-2021>.




Figure 1: Quantile regret to the top K/n proportion of experts versus number of experts N .

E.1 Quantile regret

For this experiment, we follow the same setup as Chaudhuri, Freund, and Hsu [CFH09]. The losses are generated from the *Hadamard* matrix of dimension 64, where the row with constant values is removed, the remaining rows are duplicated with their signs inverted, and all the rows are repeated horizontally in order to get $T = 32768$ columns and $n = 126$ rows. Each row i represents the losses for expert i from $t = 1, \dots, T$. Next, given a parameter K , which represents the number of *good* experts (out of n), the value 0.025 is subtracted from each entry in the first K rows. Each row is then duplicated N/n times in order to have N total experts. Finally, the values in the matrix are shifted and normalized in order to be bounded in the range $[0, 1]$. For more details, see Chaudhuri, Freund, and Hsu [CFH09, Section 3]. As noted by Chaudhuri, Freund, and Hsu [CFH09], the replication factor affects the behaviour of algorithms that are tuned in terms of the total number of experts N , such as Hedge with a decreasing regularizer scaling [MG19] or AdaHedge [d+14]. On the other hand, parameter-free algorithms such as Squint [Kv15], NormalHedge [CFH09] or CoinBetting [OP16] are not affected by the replication process. The results illustrated in Fig. 1 show that abNormal follows the same behaviour as parameter-free algorithms. In particular, we plot the cumulative regret after $T = 32768$ time-steps against the best expert. Note that the cumulative loss is equal for the $K \times N/n$ *good* experts, so we are effectively plotting the regret against the top $\varepsilon = K/n$ -quantile, which is kept constant despite the total number of experts N increasing. We adopt an initial uniform distribution over experts for all the algorithms.

E.2 Semi-adversarial losses

For this experiment, we compare Hedge with decreasing regularizer scaling $\eta_t \propto 1/\sqrt{t}$, Meta-CARE [BNR20] and FTRL-CARL from Section 4. In this case, we keep the number of total experts fixed to $N = 1000$, and use deterministic losses. When $N_0 = 1$, the best expert has loss 0.4 and the rest have loss 0.5 on every round. When $N_0 = 2$, expert 1 alternates between loss of 0 and loss of 1, expert 2 alternates (on opposite rounds from expert 1) between loss of 1 and loss of 0, and the rest have loss 0.6 on every round. Thus, both $N_0 = 1$ and $N_0 = 2$ correspond to $\Delta_0 = 0.1$. When $N_0 = N$, the first half and second half of the experts alternate between loss 1 and loss 0 on opposite rounds.

As the theory of Bilodeau, Negrea, and Roy [BNR20] prescribes, Hedge does well for $N_0 = 1$ but does equally poorly for $N_0 = 2$ as $N_0 = N$, while Meta-CARE improves on Hedge for $1 < N_0 < N$. As Theorem 3 prescribes, FTRL-CARL does well in all settings, improving on Meta-CARE by adapting faster in the $N_0 = 2$ case.

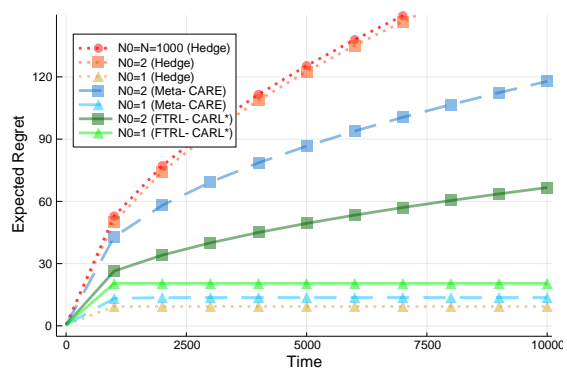


Figure 2: Semi-adversarial paradigm: Expected regret versus time for various values of N_0 .