



HAL
open science

Online non-convex optimization with imperfect feedback

Amélie Héliou, Matthieu Martin, Panayotis Mertikopoulos, Thibaud Rahier

► To cite this version:

Amélie Héliou, Matthieu Martin, Panayotis Mertikopoulos, Thibaud Rahier. Online non-convex optimization with imperfect feedback. NeurIPS 2020 - 34th International Conference on Neural Information Processing Systems, 2020, Vancouver, Canada. ⟨hal-03043746⟩

HAL Id: hal-03043746

<https://inria.hal.science/hal-03043746v1>

Submitted on 7 Dec 2020

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

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



HAL Authorization

ONLINE NON-CONVEX OPTIMIZATION WITH IMPERFECT FEEDBACK

AMÉLIE HÉLIOU*, MATTHIEU MARTIN*,
PANAYOTIS MERTIKOPOULOS^{◇,*}, AND THIBAUD RAHIER*

ABSTRACT. We consider the problem of online learning with non-convex losses. In terms of feedback, we assume that the learner observes – or otherwise constructs – an inexact model for the loss function encountered at each stage, and we propose a mixed-strategy learning policy based on dual averaging. In this general context, we derive a series of tight regret minimization guarantees, both for the learner’s static (external) regret, as well as the regret incurred against the best *dynamic* policy in hindsight. Subsequently, we apply this general template to the case where the learner only has access to the actual loss incurred at each stage of the process. This is achieved by means of a kernel-based estimator which generates an inexact model for each round’s loss function using only the learner’s realized losses as input.

1. INTRODUCTION

In this paper, we consider the following online learning framework:

- (1) At each stage $t = 1, 2, \dots$ of a repeated decision process, the learner selects an action x_t from a compact convex subset \mathcal{K} of a Euclidean space \mathbb{R}^n .
- (2) The agent’s choice of action triggers a loss $\ell_t(x_t)$ based on an a priori unknown *loss function* $\ell_t: \mathcal{K} \rightarrow \mathbb{R}$; subsequently, the process repeats.

If the loss functions ℓ_t encountered by the agent are *convex*, the above framework is the standard online convex optimization setting of Zinkevich [58] – for a survey, see [16, 29, 45] and references therein. In this case, simple first-order methods like online gradient descent (OGD) allow the learner to achieve $\mathcal{O}(T^{1/2})$ regret after T rounds [58], a bound which is well-known to be min-max optimal in this setting [1, 45]. At the same time, it is also possible to achieve tight regret minimization guarantees against *dynamic comparators* – such as the regret incurred against the best *dynamic policy* in hindsight, cf. [11, 21, 23, 30, 33] and references therein.

On the other hand, when the problem’s loss functions are not convex, the situation is considerably more difficult. When the losses are generated from a stationary stochastic distribution, the problem can be seen as a version of a continuous-armed bandit in the spirit of Agrawal [3]; in this case, there exist efficient algorithms guaranteeing logarithmic regret by discretizing the problem’s search domain and using a UCB-type policy [18, 37, 48]. Otherwise, in an adversarial context, an informed adversary can impose *linear* regret to

* CRITEO AI LAB, FRANCE

◇ UNIV. GRENoble ALPES, CNRS, INRIA, LIG, 38000, GRENoble, FRANCE.

E-mail addresses: a.heliou@criteo.com, mat.martin@criteo.com, panayotis.mertikopoulos@imag.fr, t.rahier@criteo.com.

2020 *Mathematics Subject Classification.* Primary 68Q32; secondary 90C26, 91A26.

Key words and phrases. Online optimization; non-convex; dual averaging; bandit / imperfect feedback.

P. Mertikopoulos is grateful for financial support by the French National Research Agency (ANR) in the framework of the “Investissements d’avenir” program (ANR-15-IDEX-02), the LabEx PERSYVAL (ANR-11-LABX-0025-01), and MIAI@Grenoble Alpes (ANR-19-P3IA-0003). This research was also supported by the COST Action CA16228 “European Network for Game Theory” (GAMENET).

any *deterministic algorithm* employed by the learner [31, 45, 50]; as a result, UCB-type approaches are no longer suitable.

In view of this impossibility result, two distinct threads of literature have emerged for online non-convex optimization. One possibility is to examine less demanding measures of regret – like the learner’s *local regret* [31] – and focus on first-order methods that minimize it efficiently [28, 31]. Another possibility is to consider *randomized* algorithms, in which case achieving no regret *is* possible: Krichene et al. [38] showed that adapting the well-known *Hedge* (or multiplicative / exponential weights) algorithm to a continuum allows the learner to achieve $\mathcal{O}(T^{1/2})$ regret, as in the convex case. This result is echoed in more recent works by Agarwal et al. [2] and Suggala & Netrapalli [50] who analyzed the “follow the perturbed leader” (FTPL) algorithm of Kalai & Vempala [34] with exponentially distributed perturbations and an offline optimization oracle (exact or approximate); again, the regret achieved by FTPL in this setting is $\mathcal{O}(T^{1/2})$, i.e., order-equivalent to that of Hedge in a continuum.

Our contributions and related work. A crucial assumption in the above works on randomized algorithms is that, after selecting an action, the learner receives perfect information on the loss function encountered – i.e., an *exact model* thereof. This is an important limitation for the applicability of these methods, which led to the following question by Krichene et al. [38, p. 8]:

One question is whether one can generalize the Hedge algorithm to a bandit setting, so that sublinear regret can be achieved without the need to explicitly maintain a cover.

To address this open question, we begin by considering a general framework for randomized action selection with *imperfect feedback* – i.e., with an inexact model of the loss functions encountered at each stage. Our contributions in this regard are as follows:

- (1) We present a flexible algorithmic template for online non-convex learning based on dual averaging with imperfect feedback [42].
- (2) We provide tight regret minimization rates – both *static* and *dynamic* – under a wide range of different assumptions for the loss models available to the optimizer.
- (3) We show how this framework can be extended to learning with *bandit feedback*, i.e., when the learner only observes their realized loss and must construct a loss model from scratch.

Viewed abstractly, the dual averaging (DA) algorithm is an “umbrella” scheme that contains Hedge as a special case for problems with a simplex-like domain. In the context of online convex optimization, the method is closely related to the well-known “follow the regularized leader” (FTRL) algorithm of Shalev-Shwartz & Singer [46], the FTPL method of Kalai & Vempala [34], “lazy” mirror descent (MD) [15, 16, 45], etc. For an appetizer to the vast literature surrounding these methods, we refer the reader to [14, 16, 42, 45, 46, 55, 57] and references therein.

In the *non-convex* setting, our regret minimization guarantees can be summarized as follows (see also Table 1 above): if the learner has access to inexact loss models that are unbiased and finite in mean square, the DA algorithm achieves in expectation a *static* regret bound of $\mathcal{O}(T^{1/2})$. Moreover, in terms of the learner’s *dynamic* regret, the algorithm enjoys a bound of $\mathcal{O}(T^{2/3}V_T^{1/3})$ where $V_T := \sum_{t=1}^T \|\ell_{t+1} - \ell_t\|_\infty$ denotes the *variation* of the loss functions encountered over the horizon of play (cf. Section 4 for the details). Importantly, both bounds are order-optimal, even in the context of online *convex* optimization, cf. [1, 11, 20].

With these general guarantees in hand, we tackle the bandit setting using a “kernel smoothing” technique in the spirit of Bubeck et al. [19]. This leads to a new algorithm, which we call *bandit dual averaging* (BDA), and which can be seen as a version of the DA method with *biased* loss models. The bias of the loss model can be controlled by tuning the “radius”

Feedback	Convex Losses				Non-Convex Losses	
	Static regret		Dynamic regret		Static regret	Dynamic regret
Exact	$\mathcal{O}(T^{1/2})$	[58]	$\mathcal{O}(T^{2/3}V_T^{1/3})$	[11]	$\mathcal{O}(T^{1/2})$	[38, 50] $\mathcal{O}(T^{2/3}V_T^{1/3})$
Unbiased	$\mathcal{O}(T^{1/2})$	[58]	$\mathcal{O}(T^{2/3}V_T^{1/3})$	[11]	$\mathcal{O}(T^{1/2})$	$\mathcal{O}(T^{2/3}V_T^{1/3})$
Bandit	$\mathcal{O}(T^{1/2})$	[17, 19]	$\mathcal{O}(T^{4/5}V_T^{1/5})$	[11]	$\mathcal{O}(T^{\frac{n+2}{n+3}})$	$\mathcal{O}(T^{\frac{n+3}{n+4}}V_T^{\frac{1}{n+4}})$

Table 1: Overview of related work. In regards to feedback, an “exact” model means that the learner acquires perfect knowledge of the encountered loss functions; “unbiased” refers to an inexact model that is only accurate on average; finally, “bandit” means that the learner records their incurred loss and has no other information. We only report here the best known bounds in the literature; all bounds derived in this paper are typeset **in bold**.

of the smoothing kernel; however, this comes at the cost of increasing the model’s variance – an incarnation of the well-known “bias-variance” trade-off. By resolving this trade-off, we are finally able to answer the question of Krichene et al. [38] in the positive: BDA enjoys an $\mathcal{O}(T^{\frac{n+2}{n+3}})$ static regret bound and an $\mathcal{O}(T^{\frac{n+3}{n+4}}V_T^{1/(n+4)})$ *dynamic* regret bound, without requiring an explicit discretization of the problem’s search space.

This should be contrasted with the case of online *convex* learning, where it is possible to achieve $\mathcal{O}(T^{3/4})$ regret through the use of simultaneous perturbation stochastic approximation (SPSA) techniques [25], or even $\mathcal{O}(T^{1/2})$ by means of kernel-based methods [17, 19]. This represents a drastic drop from $\mathcal{O}(T^{1/2})$, but this cannot be avoided: the worst-case bound for stochastic non-convex optimization is $\Omega(T^{(n+1)/(n+2)})$ [36, 37], so our static regret bound is nearly optimal in this regard (i.e., up to $\mathcal{O}(T^{-1/(n+2)(n+3)})$, a term which is insignificant for horizons $T \leq 10^{12}$). Correspondingly, in the case of dynamic regret minimization, the best known upper bound is $\mathcal{O}(T^{4/5}V_T^{1/5})$ for online *convex* problems [11, 24]. We are likewise not aware of any comparable dynamic regret bounds for online *non-convex* problems; to the best of our knowledge, our paper is the first to derive dynamic regret guarantees for online non-convex learning with bandit feedback.

We should stress here that, as is often the case for methods based on lifting, much of the computational cost is hidden in the sampling step. This is also the case for the proposed DA method which, like [38], implicitly assumes access to a sampling oracle. Estimating (and minimizing) the per-iteration cost of sampling is an important research direction, but one that lies beyond the scope of the current paper, so we do not address it here.

2. SETUP AND PRELIMINARIES

2.1. The model. Throughout the sequel, our only blanket assumption will be as follows:

Assumption 1. The stream of loss functions encountered is *uniformly bounded Lipschitz*, i.e., there exist constants $R, L > 0$ such that:

- (1) $|\ell_t(x)| \leq R$ for all $x \in \mathcal{K}$; more succinctly, $\|\ell_t\|_\infty \leq R$.
- (2) $|\ell_t(x') - \ell_t(x)| \leq L\|x' - x\|$ for all $x, x' \in \mathcal{K}$.

Other than this meager regularity requirement, we make no structural assumptions for ℓ_t (such as convexity, unimodality, or otherwise). In this light, the framework under consideration is akin to the online non-convex setting of Krichene et al. [38], Hazan et al. [31], and Suggala & Netrapalli [50]. The main difference with the setting of Krichene et al. [38] is that the problem’s domain \mathcal{K} is assumed convex; this is done for convenience only, to avoid technical subtleties involving “uniform fatness” conditions and the like.

In terms of playing the game, we will assume that the learner can employ *mixed strategies* to randomize their choice of action at each stage; however, because this mixing occurs over a *continuous* domain, defining this randomization requires some care. To that end, let $\mathcal{M} \equiv \mathcal{M}(\mathcal{K})$ denote the space of all finite signed Radon measures on \mathcal{K} . Then, a *mixed strategy* is defined as an element π of the set of Radon probability measures $\Delta \equiv \Delta(\mathcal{K}) \subseteq \mathcal{M}(\mathcal{K})$ on \mathcal{K} , and the player’s expected loss under π when facing a bounded loss function $\ell \in \mathcal{L}^\infty(\mathcal{K})$ will be denoted as

$$\langle \ell, \pi \rangle := \mathbb{E}_\pi[\ell] = \int_{\mathcal{K}} \ell(x) d\pi(x). \quad (1)$$

Remark 1. We should note here that Δ contains a vast array of strategies, including atomic and singular distributions that do not admit a density. For this reason, we will write Δ_{cont} for the set of strategies that are *absolutely continuous* relative to the Lebesgue measure λ on \mathcal{K} , and Δ_\perp for the set of singular strategies (which are not); by Lebesgue’s decomposition theorem [26], we have $\Delta = \Delta_{\text{cont}} \cup \Delta_\perp$. By construction, Δ_\perp contains the player’s *pure strategies*, i.e., Dirac point masses δ_x that select $x \in \mathcal{K}$ with probability 1; however, it also contains pathological strategies that admit *neither* a density, *nor* a point mass function – such as the Cantor distribution [26]. By contrast, the Radon-Nikodym (RN) derivative $p := d\pi/d\lambda$ of π exists for all $\pi \in \Delta_{\text{cont}}$, so we will sometimes refer to elements of Δ_{cont} as “Radon-Nikodym strategies”; in particular, if $\pi \in \Delta_{\text{cont}}$, we will not distinguish between π and p unless absolutely necessary to avoid confusion.

Much of our analysis will focus on strategies χ with a piecewise constant density on \mathcal{K} , i.e., $\chi = \sum_{i=1}^k \alpha_i \mathbb{1}_{\mathcal{C}_i}$ for a collection of weights $\alpha_i \geq 0$ and measurable subsets $\mathcal{C}_i \subseteq \mathcal{K}$, $i = 1, \dots, k$, such that $\int_{\mathcal{K}} \chi = \sum_i \alpha_i \lambda(\mathcal{C}_i) = 1$. These strategies will be called *simple* and the space of simple strategies on \mathcal{K} will be denoted by $\mathcal{X} \equiv \mathcal{X}(\mathcal{K})$. A key fact regarding simple strategies is that \mathcal{X} is dense in Δ in the weak topology of \mathcal{M} [26, Chap. 3]; as a result, the learner’s expected loss under *any* mixed strategy $\pi \in \Delta$ can be approximated within arbitrary accuracy $\varepsilon > 0$ by a simple strategy $\chi \in \mathcal{X}$. In addition, when k (or n) is not too large, sampling from simple strategies can be done efficiently; for all these reasons, simple strategies will play a key role in the sequel.

2.2. Measures of regret. With all this in hand, the *regret* of a learning policy $\pi_t \in \Delta$, $t = 1, 2, \dots$, against a benchmark strategy $\pi^* \in \Delta$ is defined as

$$\text{Reg}_{\pi^*}(T) = \sum_{t=1}^T [\mathbb{E}_{\pi_t}[\ell_t] - \mathbb{E}_{\pi^*}[\ell_t]] = \sum_{t=1}^T \langle \ell_t, \pi_t - \pi^* \rangle, \quad (2)$$

i.e., as the difference between the player’s mean cumulative loss under π_t and π^* over T rounds. In a slight abuse of notation, we write $\text{Reg}_{p^*}(T)$ if π^* admits a density p^* , and $\text{Reg}_x(T)$ for the regret incurred against the pure strategy δ_x , $x \in \mathcal{K}$. Then, the player’s (*static*) *regret* under π_t is given by

$$\text{Reg}(T) = \max_{x \in \mathcal{K}} \text{Reg}_x(T) = \sup_{\chi \in \mathcal{X}} \text{Reg}_\chi(T) \quad (3)$$

where the maximum is justified by the compactness of \mathcal{K} and the continuity of each ℓ_t . The lemma below provides a link between pure comparators and their approximants in the spirit of [Krichene et al. \[38\]](#); to streamline our discussion, we defer the proof to the supplement:

Lemma 1. *Let \mathcal{U} be a convex neighborhood of x in \mathcal{K} and let $\chi \in \mathcal{X}$ be a simple strategy supported on \mathcal{U} . Then, $\text{Reg}_x(T) \leq \text{Reg}_\chi(T) + L \text{diam}(\mathcal{U})T$.*

This lemma will be used to bound the agent’s static regret using bounds obtained for simple strategies $\chi \in \mathcal{X}$. Going beyond static comparisons of this sort, the learner’s *dynamic regret* is defined as

$$\text{DynReg}(T) = \sum_{t=1}^T [\langle \ell_t, \pi_t \rangle - \min_{\pi \in \Delta} \langle \ell_t, \pi \rangle] = \sum_{t=1}^T \langle \ell_t, \pi_t - \pi_t^* \rangle \quad (4)$$

where $\pi_t^* \in \arg \min_{\pi \in \Delta} \langle \ell_t, \pi \rangle$ is a “best-response” to ℓ_t (that such a strategy exists is a consequence of the compactness of \mathcal{K} and the continuity of each ℓ_t). In regard to its static counterpart, the agent’s dynamic regret is considerably more ambitious, and achieving sublinear dynamic regret is not always possible; we examine this issue in detail in [Section 4](#).

2.3. Feedback models. After choosing an action, the agent is only assumed to observe an *inexact model* $\hat{u}_t \in \mathcal{L}^\infty(\mathcal{K})$ of the t -th stage loss function ℓ_t ; for concreteness, we will write

$$\hat{u}_t = \ell_t + e_t \quad (5)$$

where the “observation error” e_t captures all sources of uncertainty in the player’s model. This uncertainty could be both “random” (zero-mean) or “systematic” (non-zero-mean), so it will be convenient to decompose e_t as

$$e_t = z_t + b_t \quad (6)$$

where z_t is zero-mean and b_t denotes the mean of e_t .

To define all this formally, we will write $\mathcal{F}_t = \mathcal{F}(\pi_1, \dots, \pi_t)$ for the history of the player’s mixed strategy up to stage t (inclusive). The chosen action x_t and the observed model \hat{u}_t are both generated *after* the player chooses π_t so, by default, they are not \mathcal{F}_t -measurable. Accordingly, we will collect all randomness affecting \hat{u}_t in an abstract probability law \mathbb{P} , and we will write $b_t = \mathbb{E}[e_t | \mathcal{F}_t]$ and $z_t = e_t - b_t$; in this way, $\mathbb{E}[z_t | \mathcal{F}_t] = 0$ by definition.

In view of all this, we will focus on the following descriptors for \hat{u}_t :

$$a) \text{ Bias: } \quad \|b_t\|_\infty \leq B_t \quad (7a)$$

$$b) \text{ Variance: } \quad \mathbb{E}[\|z_t\|_\infty^2 | \mathcal{F}_t] \leq \sigma_t^2 \quad (7b)$$

$$c) \text{ Mean square: } \quad \mathbb{E}[\|\hat{u}_t\|_\infty^2 | \mathcal{F}_t] \leq M_t^2 \quad (7c)$$

In the above, B_t , σ_t and M_t are deterministic constants that are to be construed as bounds on the bias, (conditional) variance, and magnitude of the model \hat{u}_t at time t . In obvious terminology, a model with $B_t = 0$ will be called *unbiased*, and an unbiased model with $\sigma_t = 0$ will be called *exact*.

► **Example 1** (Parametric models). An important application of online optimization is the case where the encountered loss functions are of the form $\ell_t(x) = \ell(x; \theta_t)$ for some sequence of parameter vectors $\theta_t \in \mathbb{R}^m$. In this case, the learner typically observes an estimate $\hat{\theta}_t$ of θ_t , leading to the inexact model $\hat{u}_t = \ell(\cdot; \hat{\theta}_t)$. Importantly, this means that \hat{u}_t *does not require infinite-dimensional feedback* to be constructed. Moreover, the dependence of ℓ on θ is often linear, so if $\hat{\theta}_t$ is an unbiased estimate of θ_t , then so is \hat{u}_t . ◀

► **Example 2** (Online clique prediction). As a specific incarnation of a parametric model, consider the problem of finding the largest complete subgraph – a *maximum clique* – of an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. This is a key problem in machine learning with applications to social networks [27], data mining [9], gene clustering [49], feature embedding [56], and many other fields. In the online version of the problem, the learner is asked to predict such a clique in a graph \mathcal{G}_t that evolves over time (e.g., a social network), based on partial historical observations of the graph. Then, by the Motzkin-Straus theorem [13, 41], this boils down to an online quadratic program of the form:

$$\text{maximize } u_t(x) = \sum_{i,j=1}^n x_i A_{ij,t} x_j \quad \text{subject to } x_i \geq 0, \sum_{i=1}^n x_i = 1, \quad (\text{MCP})$$

where $A_t = (A_{ij,t})_{i,j=1}^n$ denotes the adjacency matrix of \mathcal{G}_t . Typically, \hat{A}_t is constructed by picking a node i uniformly at random, charting out its neighbors, and letting $\hat{A}_{ij,t} = |\mathcal{V}|/2$

whenever j is connected to i . It is easy to check that \hat{A}_t is an unbiased estimator of A_t ; as a result, the function $\hat{u}_t(x) = x^\top \hat{A}_t x$ is an unbiased model of u_t . ◀

► **Example 3** (Online-to-batch). Consider an empirical risk minimization model of the form

$$f(x) = \frac{1}{m} \sum_{i=1}^m f_i(x) \quad (8)$$

where each $f_i: \mathbb{R}^n \rightarrow \mathbb{R}$ corresponds to a data point (or “sample”). In the “online-to-batch” formulation of the problem [45], the optimizer draws uniformly at random a sample $i_t \in \{1, \dots, m\}$ at each stage $t = 1, 2, \dots$, and observes $\hat{\ell}_t = f_{i_t}$. Typically, each f_i is relatively easy to store in closed form, so $\hat{\ell}_t$ is an easily available unbiased model of the empirical risk function f . ◀

3. PROX-STRATEGIES AND DUAL AVERAGING

The class of non-convex online learning policies that we will consider is based on the general template of *dual averaging* (DA) / “follow the regularized leader” (FTRL) methods. Informally, this scheme can be described as follows: at each stage $t = 1, 2, \dots$, the learner plays a mixed strategy that minimizes their cumulative loss up to round $t - 1$ (inclusive) plus a “regularization” penalty term (hence the “regularized leader” terminology). In the rest of this section, we provide a detailed construction and description of the method.

3.1. Randomizing over discrete vs. continuous sets. We begin by describing the dual averaging method when the underlying action set is *finite*, i.e., of the form $\mathcal{A} = \{1, \dots, n\}$. In this case, the space of mixed strategies is the n -dimensional simplex $\Delta_n := \Delta(\mathcal{A}) = \{p \in \mathbb{R}_+^n : \sum_{i=1}^n p_i = 1\}$, and, at each $t = 1, 2, \dots$, the dual averaging algorithm prescribes the mixed strategy

$$p_t \leftarrow \arg \min_{p \in \Delta_n} \{\eta \sum_{s=1}^{t-1} \langle \hat{u}_s, p \rangle + h(p)\}. \quad (9)$$

In the above, $\eta > 0$ is a “learning rate” parameter and $h: \Delta_n \rightarrow \mathbb{R}$ is the method’s “regularizer”, assumed to be continuous and strongly convex over Δ_n . In this way, the algorithm can be seen as tracking the “best” choice up to the present, modulo a “day 0” regularization component – the “follow the regularized leader” interpretation.

In our case however, the method is to be applied to the *infinite-dimensional* set $\Delta \equiv \Delta(\mathcal{K})$ of the learner’s mixed strategies, so the issue becomes considerably more involved. To illustrate the problem, consider one of the prototypical regularizer functions, the negentropy $h(p) = \sum_{i=1}^n p_i \log p_i$ on Δ_n . If we naïvely try to extend this definition to the infinite-dimensional space $\Delta(\mathcal{K})$, we immediately run into problems: First, for pure strategies, any expression of the form $\sum_{x' \in \mathcal{K}} \delta_x(x') \log \delta_x(x')$ would be meaningless. Second, even if we focus on Radon-Nikodym strategies $p \in \Delta_{\text{cont}}$ and use the integral definition $h(p) = \int_{\mathcal{K}} p \log p$, a density like $p(x) \propto 1/(x \log x)^2$ on $\mathcal{K} = [0, 1/2]$ has *infinite* negentropy, implying that even Δ_{cont} is too large to serve as a domain.

3.2. Formal construction of the algorithm. To overcome the issues identified above, our starting point will be that any mixed-strategy incarnation of the dual averaging algorithm must contain *at least* the space $\mathcal{X} \equiv \mathcal{X}(\mathcal{K})$ of the player’s simple strategies. To that end, let \mathcal{V} be an ambient Banach space which contains the set of simple strategies \mathcal{X} as an embedded subset. For technical reasons, we will also assume that the topology induced on \mathcal{X} by the reference norm $\|\cdot\|$ of \mathcal{V} is not weaker than the natural topology on \mathcal{X} induced by the total variation norm; formally, $\|\cdot\|_{\text{TV}} \leq \alpha \|\cdot\|$ for some $\alpha > 0$.¹ For example, \mathcal{V} could be the (Banach) space $\mathcal{M}(\mathcal{K})$ of finite signed measures on \mathcal{K} , the (Hilbert) space $\mathcal{L}^2(\mathcal{K})$ of square

¹Since the dual space of $\mathcal{M}(\mathcal{K})$ contains $\mathcal{L}^\infty(\mathcal{K})$, we will also view $\mathcal{L}^\infty(\mathcal{K})$ as an embedded subset of \mathcal{V}^* .

integrable functions on \mathcal{K} endowed with the L^2 norm,² or an altogether different model for \mathcal{X} . We then have:

Definition 1. A *regularizer* on \mathcal{X} is a lower semi-continuous (l.s.c.) convex function $h: \mathcal{V} \rightarrow \mathbb{R} \cup \{\infty\}$ such that:

- (1) \mathcal{X} is a weakly dense subset of the *effective domain* $\text{dom } h := \{p: h(p) < \infty\}$ of h .
- (2) The subdifferential ∂h of h admits a *continuous selection*, i.e., there exists a continuous mapping ∇h on $\text{dom } \partial h := \{\partial h \neq \emptyset\}$ such that $\nabla h(q) \in \partial h(q)$ for all $q \in \text{dom } \partial h$.
- (3) h is strongly convex, i.e., there exists some $K > 0$ such that $h(p) \geq h(q) + \langle \nabla h(q), p - q \rangle + (K/2)\|p - q\|^2$ for all $p \in \text{dom } h, q \in \text{dom } \partial h$.

The set $\mathcal{Q} := \text{dom } \partial h$ will be called the *prox-domain* of h ; its elements will be called *prox-strategies*.

Remark. For completeness, recall that the subdifferential of h at q is the set $\partial h(q) = \{\psi \in \mathcal{V}^*: h(p) \geq h(q) + \langle \psi, p - q \rangle \text{ for all } p \in \mathcal{V}\}$; also lower semicontinuity means that the sublevel sets $\{h \leq c\}$ of h are closed for all $c \in \mathbb{R}$. For more details, we refer the reader to Phelps [44].

Some prototypical examples of this general framework are as follows (with more in the supplement):

► **Example 4** (L^2 regularization). Let $\mathcal{V} = \mathcal{L}^2(\mathcal{K})$ and consider the quadratic regularizer $h(p) = (1/2)\|p\|_2^2 = (1/2) \int_{\mathcal{K}} p^2$ if $p \in \Delta_{\text{cont}} \cap \mathcal{L}^2(\mathcal{K})$, and $h(p) = \infty$ otherwise. In this case, $\mathcal{Q} = \text{dom } h = \Delta_{\text{cont}} \cap \mathcal{L}^2(\mathcal{K})$ and $\nabla h(q) = q$ is a continuous selection of ∂h on \mathcal{Q} . ◀

► **Example 5** (Entropic regularization). Let $\mathcal{V} = \mathcal{M}(\mathcal{K})$ and consider the entropic regularizer $h(p) = \int_{\mathcal{K}} p \log p$ whenever p is a density with finite entropy, $h(p) = \infty$ otherwise. By Pinsker's inequality, h is 1-strongly convex relative to the total variation norm $\|\cdot\|_{\text{TV}}$ on \mathcal{V} ; moreover, we have $\mathcal{Q} = \{q \in \Delta_{\text{cont}} : \text{supp}(q) = \mathcal{K}\} \subsetneq \text{dom } h$ and $\nabla h(q) = 1 + q \log q$ on \mathcal{Q} . In the finite-dimensional case, this regularizer forms the basis of the well-known Hedge (or multiplicative/exponential weights) algorithm [5, 6, 39, 54]; for the infinite-dimensional case, see [38, 43] (and below). ◀

With all this in hand, the dual averaging algorithm can be described by means of the abstract recursion

$$y_{t+1} = y_t - \hat{u}_t, \quad p_{t+1} = Q(\eta_{t+1} y_{t+1}), \quad (\text{DA})$$

where (i) $t = 1, 2, \dots$ denotes the stage of the process (with the convention $y_0 = \hat{u}_0 = 0$); (ii) $p_t \in \mathcal{Q}$ is the learner's strategy at stage t ; (iii) $\hat{u}_t \in \mathcal{L}^\infty(\mathcal{K})$ is the inexact model revealed at stage t ; (iv) $y_t \in \mathcal{L}^\infty(\mathcal{K})$ is a "score" variable that aggregates loss models up to stage t ; (v) $\eta_t > 0$ is a "learning rate" sequence; and (vi) $Q: \mathcal{L}^\infty(\mathcal{K}) \rightarrow \mathcal{Q}$ is the method's *mirror map*, viz.

$$Q(\psi) = \arg \max_{p \in \mathcal{V}} \{\langle \psi, p \rangle - h(p)\}. \quad (10)$$

For a pseudocode implementation, see Alg. 1 above. In the paper's supplement we also show that the method is *well-posed*, i.e., the arg max in (10) is attained at a valid prox-strategy $p_t \in \mathcal{Q}$. We illustrate this with an example:

► **Example 6** (Logit choice). Suppose that $h(p) = \int_{\mathcal{K}} p \log p$ is the entropic regularizer of Example 5. Then, the corresponding mirror map is given in closed form by the logit choice model:

$$\Lambda(\varphi) = \frac{\exp(\varphi)}{\int_{\mathcal{K}} \exp(\varphi)} \quad \text{for all } \varphi \in \mathcal{L}^\infty(\mathcal{K}). \quad (11)$$

²In this case, $\alpha = \sqrt{\lambda(\mathcal{K})}$: this is because $\|p\|_{\text{TV}}^2 = [\int_{\mathcal{K}} p]^2 \leq \int_{\mathcal{K}} p^2 \cdot \int_{\mathcal{K}} 1 = \lambda(\mathcal{K}) \|p\|_2^2$ if $p \in \Delta_{\text{cont}}$.

Algorithm 1: Dual averaging with imperfect feedback [Hedge variant: $Q \leftarrow \Lambda$]

Require: mirror map $Q: \mathcal{L}^\infty(\mathcal{K}) \rightarrow \mathcal{Q}$; learning rate $\eta_t > 0$; **initialize:** $y_1 \leftarrow 0$

```

1: for  $t = 1, 2, \dots$  do
2:   set  $p_t \leftarrow Q(\eta_t y_t)$       [ $p_t \leftarrow \Lambda(\eta_t y_t)$  for Hedge]      # update mixed strategy
3:   play  $x_t \sim p_t$                                      # choose action
4:   observe  $\hat{u}_t$                                        # model revealed
5:   set  $y_{t+1} \leftarrow y_t - \hat{u}_t$                        # update scores
6: end for

```

This derivation builds on a series of well-established arguments that we defer to the supplement. Clearly, $\int_{\mathcal{K}} \Lambda(\varphi) = 1$ and $\Lambda(\varphi) > 0$ as a function on \mathcal{K} , so $\Lambda(\varphi)$ is a valid prox-strategy. \blacktriangleleft

4. GENERAL REGRET BOUNDS

4.1. Static regret guarantees. We are now in a position to state our first result for (DA):

Proposition 1. *For any simple strategy $\chi \in \mathcal{X}$, Alg. 1 enjoys the bound*

$$\text{Reg}_\chi(T) \leq \eta_{T+1}^{-1} [h(\chi) - \min h] + \sum_{t=1}^T \langle e_t, \chi - p_t \rangle + \frac{1}{2K} \sum_{t=1}^T \eta_t \|\hat{u}_t\|_*^2. \quad (12)$$

Proposition 1 is a “template” bound that we will use to extract static and dynamic regret guarantees in the sequel. Its proof relies on the introduction of a suitable energy function measuring the match between the learner’s aggregate model y_t and the comparator χ . The main difficulty is that these variables live in completely different spaces ($\mathcal{L}^\infty(\mathcal{K})$ vs. \mathcal{X} respectively), so there is no clear distance metric connecting them. However, since bounded functions $\psi \in \mathcal{L}^\infty(\mathcal{K})$ and simple strategies $\chi \in \mathcal{X}$ are naturally paired via duality, they are indirectly connected via the Fenchel–Young inequality $\langle \psi, \chi \rangle \leq h(\chi) + h^*(\psi)$, where $h^*(\psi) = \max_{p \in \mathcal{V}} \{\langle \psi, p \rangle - h(p)\}$ denotes the convex conjugate of h and equality holds if and only if $Q(\psi) = \chi$. We will thus consider the energy function

$$E_t := \eta_t^{-1} [h(\chi) + h^*(\eta_t y_t) - \langle \eta_t y_t, \chi \rangle]. \quad (13)$$

By construction, $E_t \geq 0$ for all t and $E_t = 0$ if and only if $p_t = Q(\eta_t y_t) = \chi$. More to the point, the defining property of E_t is the following recursive bound (which we prove in the supplement):

Lemma 2. *For all $\chi \in \mathcal{X}$, we have:*

$$E_{t+1} \leq E_t + \langle \hat{u}_t, \chi - p_t \rangle + (\eta_{t+1}^{-1} - \eta_t^{-1}) [h(\chi) - \min h] + \frac{\eta_t}{2K} \|\hat{u}_t\|_*^2. \quad (14)$$

Proposition 1 is obtained by telescoping (14); subsequently, to obtain a regret bound for Alg. 1, we must relate $\text{Reg}_x(T)$ to $\text{Reg}_\chi(T)$. This can be achieved by invoking Lemma 1 but the resulting expressions are much simpler when h is *decomposable*, i.e., $h(p) = \int_{\mathcal{K}} \theta(p(x)) dx$ for some C^2 function $\theta: [0, \infty) \rightarrow \mathbb{R}$ with $\theta'' > 0$. In this more explicit setting, we have:

Theorem 1. *Fix $x \in \mathcal{K}$, let \mathcal{C} be a convex neighborhood of x in \mathcal{K} , and suppose that Alg. 1 is run with a decomposable regularizer $h(p) = \int_{\mathcal{K}} \theta \circ p$. Then, letting $\phi(z) = z\theta(1/z)$ for $z > 0$, we have:*

$$\mathbb{E}[\text{Reg}_x(T)] \leq \frac{\phi(\lambda(\mathcal{C})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + L \text{diam}(\mathcal{C})T + 2 \sum_{t=1}^T B_t + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t M_t^2. \quad (15)$$

In particular, if Alg. 1 is run with learning rate $\eta_t \propto 1/t^\rho$, $\rho \in (0, 1)$, and inexact models such that $B_t = \mathcal{O}(1/t^\beta)$ and $M_t^2 = \mathcal{O}(t^{2\mu})$ for some $\beta, \mu \geq 0$, we have:

$$\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(\phi(T^{-n\kappa})T^\rho + T^{1-\kappa} + T^{1-\beta} + T^{1+2\mu-\rho}) \quad \text{for all } \kappa \geq 0. \quad (16)$$

Corollary 1. *If the learner’s feedback is unbiased and bounded in mean square (i.e., $B_t = 0$ and $\sup_t M_t < \infty$), running Alg. 1 with learning rate $\eta_t \propto 1/t^\rho$ guarantees*

$$\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(\phi(T^{-n\rho})T^\rho + T^{1-\rho}). \quad (17)$$

In particular, for the regularizers of Examples 4 and 5, we have:

- (1) *For $\theta(z) = (1/2)z^2$, Alg. 1 with $\eta_t \propto t^{-1/(n+2)}$ guarantees $\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^{\frac{n+1}{n+2}})$.*
- (2) *For $\theta(z) = z \log z$, Alg. 1 with $\eta_t \propto t^{-1/2}$ guarantees $\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^{1/2})$.*

Remark 2. Here and in the sequel, logarithmic factors are ignored in the Landau $\mathcal{O}(\cdot)$ notation. We should also stress that the role of \mathcal{C} in Theorem 1 only has to do with the analysis of the algorithm, not with the derived bounds (which are obtained by picking a suitable \mathcal{C}).

First, in online *convex* optimization, dual averaging with stochastic gradient feedback achieves $\mathcal{O}(\sqrt{T})$ regret *irrespective* of the choice of regularizer, and this bound is tight [1, 16, 45]. By contrast, in the non-convex setting, the choice of regularizer has a visible impact on the regret because it affects the exponent of T : in particular, L^2 regularization carries a much worse dependence on T relative to the Hedge variant of Alg. 1. This is due to the term $\mathcal{O}(\phi(T^{-n\kappa})T^\rho)$ that appears in (16) and is in turn linked to the choice of the “enclosure set” \mathcal{C} having $\lambda(\mathcal{C}) \propto T^{-n\kappa}$ for some $\kappa \geq 0$.

The negentropy regularizer (and any other regularizer with quasi-linear growth at infinity, see the supplement for additional examples) only incurs a logarithmic dependence on $\lambda(\mathcal{C})$. Instead, the quadratic growth of the L^2 regularizer induces an $\mathcal{O}(1/\lambda(\mathcal{C}))$ term in the algorithm’s regret, which is ultimately responsible for the catastrophic dependence on the dimension of \mathcal{K} . Seeing as the bounds achieved by the Hedge variant of Alg. 1 are optimal in this regard, we will concentrate on this specific instance in the sequel.

4.2. Dynamic regret guarantees. We now turn to the dynamic regret minimization guarantees of Alg. 1. In this regard, we note first that, in complete generality, dynamic regret minimization is *not possible* because an informed adversary can always impose a uniformly positive loss at each stage [45]. Because of this, dynamic regret guarantees are often stated in terms of the *variation* of the loss functions encountered, namely

$$V_T := \sum_{t=1}^T \|\ell_{t+1} - \ell_t\|_\infty = \sum_{t=1}^T \max_{x \in \mathcal{K}} |\ell_{t+1}(x) - \ell_t(x)|, \quad (18)$$

with the convention $\ell_{t+1} = \ell_t$ for $t = T$.³ We then have:

Theorem 2. *Suppose that the Hedge variant of Alg. 1 is run with learning rate $\eta_t \propto 1/t^\rho$ and inexact models with $B_t = \mathcal{O}(1/t^\beta)$ and $M_t^2 = \mathcal{O}(t^{2\mu})$ for some $\beta, \mu \geq 0$. Then:*

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1+2\mu-\rho} + T^{1-\beta} + T^{2\rho-2\mu}V_T). \quad (19)$$

In particular, if $V_T = \mathcal{O}(T^\nu)$ for some $\nu < 1$ and the learner’s feedback is unbiased and bounded in mean square (i.e., $B_t = 0$ and $\sup_t M_t < \infty$), the choice $\rho = (1-\nu)/3$ guarantees

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{\frac{2+\nu}{3}}). \quad (20)$$

To the best of our knowledge, Theorem 2 provides the first dynamic regret guarantee for online non-convex problems. The main idea behind its proof is to examine the evolution of play over a series of windows of length $\Delta = \mathcal{O}(T^\gamma)$ for some $\gamma > 0$. In so doing, Theorem 1 can be used to obtain a bound for the learner’s regret relative to the best action $x \in \mathcal{K}$ *within each window*. Obviously, if the length of the window is chosen sufficiently small, aggregating

³This notion is due to Besbes et al. [11]. Other notions of variation have also been considered [11, 21, 23], as well as other measures of regret, cf. [30, 32]; for a survey, see [20].

the learner’s regret *per window* will be a reasonable approximation of the learner’s dynamic regret. At the same time, if the window is taken too small, the number of such windows required to cover T will be $\Theta(T)$, so this approximation becomes meaningless. As a result, to obtain a meaningful regret bound, this window-by-window examination of the algorithm must be carefully aligned with the variation V_T of the loss functions encountered by the learner. Albeit intuitive, the details required to make this argument precise are fairly subtle, so we relegate the proof of [Theorem 2](#) to the paper’s supplement.

We should also observe here that the $\mathcal{O}(T^{\frac{2+\nu}{3}})$ bound of [Theorem 2](#) is, in general, unimprovable, even if the losses are *linear*. Specifically, Besbes et al. [11] showed that, if the learner is facing a stream of linear losses with stochastic gradient feedback (i.e., an inexact linear model), an informed adversary can still impose $\text{DynReg}(T) = \Omega(T^{2/3}V_T^{1/3})$. Besbes et al. [11] further proposed a scheme to achieve this bound by means of a periodic restart meta-principle that partitions the horizon of play into batches of size $(T/V_T)^{2/3}$ and then runs an algorithm achieving $(T/V_T)^{1/3}$ regret per batch. [Theorem 2](#) differs from the results of Besbes et al. [11] in two key aspects: (a) [Alg. 1](#) does not require a periodic restart schedule (so the learner does not forget the information accrued up to a given stage); and (b) more importantly, it applies to *general* online optimization problems, without a convex structure or any other structural assumptions (though with a different feedback structure).

5. APPLICATIONS TO ONLINE NON-CONVEX LEARNING WITH BANDIT FEEDBACK

As an application of the inexact model framework of the previous sections, we proceed to consider the case where the learner only observes their realized reward $\ell_t(x_t)$ and has *no other information*. In this “bandit setting”, an inexact model is not available and must instead be constructed on the fly.

When \mathcal{K} is a finite set, ℓ_t is a $|\mathcal{K}|$ -dimensional vector, and an unbiased estimator for ℓ_t can be constructed by setting $\hat{u}_t(x) = [\mathbb{1}\{x = x_t\} / \mathbb{P}(x = x_t)] \ell_t(x_t)$ for all $x \in \mathcal{K}$. This “importance weighted” estimator is the basis for the EXP3 variant of the Hedge algorithm which is known to achieve $\mathcal{O}(T^{1/2})$ regret [8]. However, in the case of *continuous* action spaces, there is a key obstacle: if the indicator $\mathbb{1}\{x = x_t\}$ is replaced by a Dirac point mass $\delta_{x_t}(x)$, the resulting loss model $\hat{u}_t \propto \delta_{x_t}$ would no longer be a function but a generalized (singular) distribution, so the dual averaging framework of [Alg. 1](#) no longer applies.

To counter this, we will take a “smoothing” approach in the spirit of [19] and consider the estimator

$$\hat{u}_t(x) = K_t(x_t, x) \cdot \ell_t(x_t) / p_t(x_t) \tag{21}$$

where $K_t: \mathcal{K} \times \mathcal{K} \rightarrow \mathbb{R}$ is a (time-varying) *smoothing kernel*, i.e., $\int_{\mathcal{K}} K_t(x, x') dx' = 1$ for all $x \in \mathcal{K}$. For concreteness (and sampling efficiency), we will assume that losses now take values in $[0, 1]$, and we will focus on simple kernels that are supported on a neighborhood $\mathcal{U}_\delta(x) = \mathbb{B}_\delta(x) \cap \mathcal{K}$ of x in \mathcal{K} and are constant therein, i.e., $K^\delta(x, x') = [\lambda(\mathcal{U}_\delta(x))]^{-1} \mathbb{1}\{\|x' - x\| \leq \delta\}$.

The “smoothing radius” δ in the definition of K^δ will play a key role in the choice of loss model being fed to [Alg. 1](#). If δ is taken too small, K^δ will approach a point mass, so it will have low estimation error but very high variance; at the other end of the spectrum, if δ is taken too large, the variance of the induced estimator will be low, but so will its accuracy. In view of this, we will consider a flexible smoothing schedule of the form $\delta_t = 1/t^\mu$ which gradually sharpens the estimator over time as more information comes in. Then, to further protect the algorithm from getting stuck in local minima, we will also incorporate in p_t an explicit exploration term of the form $\varepsilon_t/\lambda(\mathcal{K})$.

Putting all this together, we obtain the *bandit dual averaging* (BDA) algorithm presented in pseudocode form as [Alg. 2](#) above. By employing a slight variation of the analysis presented

Algorithm 2: Bandit dual averaging	[Hedge variant: $Q \leftarrow \Lambda$]
Require: mirror map $Q: \mathcal{L}^\infty(\mathcal{K}) \rightarrow \mathcal{Q}$; parameters $\eta_t, \delta_t, \varepsilon_t > 0$; initialize: $y_1 \leftarrow 0$	
1: for $t = 1, 2, \dots$ do	
2: set $p_t \leftarrow (1 - \varepsilon_t)Q(\eta_t y_t) + \varepsilon_t/\lambda(\mathcal{K})$ [$Q \leftarrow \Lambda$ for Hedge]	# mixed strategy
3: play $x_t \sim p_t$	# choose action
4: set $\hat{u}_t = K^{\delta_t}(x_t, \cdot) \cdot \ell_t(x_t)/p_t(x_t)$	# payoff model
5: set $y_{t+1} \leftarrow y_t - \hat{u}_t$	# update scores
6: end for	

in Section 4 (basically amounting to a tighter bound in Lemma 2), we obtain the following guarantees for Alg. 2:

Proposition 2. *Suppose that the Hedge variant of Alg. 2 is run with learning rate $\eta_t \propto 1/t^\rho$ and smoothing/exploration schedules $\delta_t \propto 1/t^\mu$, $\varepsilon_t \propto 1/t^\beta$ respectively. Then, the learner enjoys the bound*

$$\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^\rho + T^{1-\mu} + T^{1-\beta} + T^{1+n\mu+\beta-\rho}). \quad (22)$$

In particular, if the algorithm is run with $\rho = (n+2)/(n+3)$ and $\mu = \beta = 1/(n+3)$, we obtain the bound $\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^{\frac{n+2}{n+3}})$.

Proposition 3. *Suppose that the Hedge variant of Alg. 2 is run with parameters as in Proposition 2 against a stream of loss functions with variation $V_T = \mathcal{O}(T^\nu)$. Then, the learner enjoys*

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1+n\mu+\beta-\rho} + T^{1-\beta} + T^{1-\mu} + T^{\nu+2\rho-n\mu-\beta}). \quad (23)$$

In particular, if the algorithm is run with $\rho = (1-\nu)(n+2)/(n+4)$ and $\mu = \beta = (1-\nu)/(n+4)$, we obtain the optimized bound $\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{\frac{n+3+\nu}{n+4}})$.

To the best of our knowledge, Proposition 3 is the first result of its kind for dynamic regret minimization in online non-convex problems with bandit feedback. We conjecture that the bounds of Propositions 2 and 3 can be tightened further to $\mathcal{O}(T^{\frac{n+1}{n+2}})$ and $\mathcal{O}(T^{\frac{n+2+\nu}{n+3}})$ by dropping the explicit exploration term; we defer this finetuning to future work.

APPENDIX A. EXAMPLES

In this appendix, we provide some more decomposable regularizers that are commonly used in the literature:

► **Example 7** (Log-barrier regularization). Let $\mathcal{V} = \mathcal{M}(\mathcal{K})$ as above and consider the so-called *Burg entropy* $h(p) = -\int_{\mathcal{K}} \log p$ [4]. In this case, $\mathcal{Q} = \text{dom } \partial h = \{q \in \Delta_{\text{cont}} : \text{supp}(q) = \mathcal{K}\} = \text{dom } h$ and $\nabla h(q) = -1/q$ on \mathcal{Q} . In the finite-dimensional case, this regularizer plays a fundamental role in the affine scaling method of Karmarkar [35], see e.g., Tseng [52], Vanderbei et al. [53] and references therein. The corresponding mirror map is obtained as follows: let $L(p; \lambda) = \int_{\mathcal{K}} \psi p + \int_{\mathcal{K}} \log p - \lambda \int_{\mathcal{K}} p$ denote the Lagrangian of the problem (10), so $q = Q(\psi)$ satisfies the first-order optimality condition

$$\psi + 1/q - \lambda = 0. \quad (\text{A.1})$$

Solving for q and integrating, we get $\int_{\mathcal{K}} (\lambda - \psi)^{-1} = \int_{\mathcal{K}} q = 1$. The function $\phi(\lambda) = \int_{\mathcal{K}} (\lambda - \psi)^{-1}$ is decreasing in λ and continuous whenever finite; moreover, since $\psi \in \mathcal{L}^\infty(\mathcal{K})$, it follows that ϕ is always finite (and hence continuous) for large enough λ , and $\lim_{\lambda \rightarrow \infty} \phi(\lambda) = 0$. Since $\sup_{\lambda} \phi(\lambda) = \infty$, there exists some maximal λ^* such that (A.1) holds (in practice, this can be located by a simple line search initialized at some $\lambda > \|\psi\|_\infty$). We thus get $Q(\psi) = (\lambda^* - \psi)^{-1}$. ◀

► **Example 8** (Tsallis entropy). A generalization of the Shannon-Gibbs entropy for nonextensive variables is the *Tsallis entropy* [51] defined here as $h(p) = \int_{\mathcal{K}} \theta(p)$ where $\theta(z) = [\gamma(1 - \gamma)]^{-1}(z - z^\gamma)$ for $\gamma \in (0, 1]$, with the continuity convention $(z - z^\gamma)/(1 - \gamma) = z \log z$ for $\gamma = 1$ (corresponding to the Shannon-Gibbs case). Working as in [Example 7](#), we have $\mathcal{Q} = \text{dom } \partial h = \{q \in \Delta_{\text{cont}} : \text{supp}(q) = \mathcal{K}\} \subsetneq \text{dom } h$, and the corresponding mirror map $q = Q(\psi)$ is obtained via the first-order stationarity equation

$$\psi - \frac{1 - \gamma q^{\gamma-1}}{\gamma(1 - \gamma)} - \lambda = 0. \quad (\text{A.2})$$

Then, solving for q yields $Q(\psi) = (1 - \gamma)^{1/(\gamma-1)} \int_{\mathcal{K}} (\mu - \psi)^{1/(\gamma-1)}$ with $\mu > \|\psi\|_\infty$ chosen so that $\int_{\mathcal{K}} Q(\psi) = 1$. ◀

APPENDIX B. BASIC PROPERTIES OF REGULARIZERS AND MIRROR MAPS

The goal of this appendix is to prove some basic results on regularizer functions and mirror maps that will be used liberally in the sequel. Versions of the results presented here already exist in the literature, but our infinite-dimensional setting introduces some subtleties that require further care. For this reason, we state and prove all required results for completeness.

We begin by recalling some definitions from the main part of the paper. First, we write $\mathcal{M} \equiv \mathcal{M}(\mathcal{K})$ for the space of all finite signed Radon measures on \mathcal{K} equipped with the total variation norm $\|\mu\|_{\text{TV}} = \mu^+(\mathcal{K}) + \mu^-(\mathcal{K})$, where μ^+ (resp. μ^-) denotes the positive (resp. negative) part of μ coming from the Hahn-Banach decomposition of signed measures on \mathcal{K} . As we discussed in [Section 3](#), we also assume given a model Banach space \mathcal{V} containing the set of simple strategies \mathcal{X} as an embedded subset and such that $\|\cdot\|_{\text{TV}} \leq \alpha \|\cdot\|$ for some $\alpha > 0$.

With all this in hand, we begin by discussing the well-posedness of [Alg. 1](#). To that end, we have the following basic result:

Lemma B.1. *Let h be a regularizer on \mathcal{X} . Then:*

(1) $Q(\psi) \in \mathcal{Q}$ for all $\psi \in \mathcal{V}^*$; in particular:

$$q = Q(\psi) \iff \psi \in \partial h(q). \quad (\text{B.1})$$

(2) If $q = Q(\psi)$ and $p \in \text{dom } h$, we have

$$\langle \nabla h(q), q - p \rangle \leq \langle \psi, q - p \rangle. \quad (\text{B.2})$$

(3) The convex conjugate $h^*(\psi) = \max_{p \in \mathcal{V}} \{\langle \psi, p \rangle - h(p)\}$ is Fréchet differentiable and satisfies

$$D_v h^*(\psi) = \langle v, Q(\psi) \rangle \quad \text{for all } \psi, v \in \mathcal{V}^*. \quad (\text{B.3})$$

Corollary 2. *Alg. 1 is well-posed, i.e., $p_t \in \mathcal{Q}$ for all $t = 1, 2, \dots$ if $\hat{u}_t \in \mathcal{L}^\infty(\mathcal{K})$.*

Proof. We proceed item by item:

(1) First, since h is strongly convex and lower semi-continuous, the maximum in (10) is attained. Hence, by Fermat's rule for subdifferentials, q solves (10) if and only if $\psi - \partial h(q) \ni 0$. We thus get the string of equivalences:

$$q = Q(\psi) \iff \psi - \partial h(q) \ni 0 \iff \psi \in \partial h(q). \quad (\text{B.4})$$

In particular, this implies that $\psi \in \text{dom } \partial h(q) \neq \emptyset$, i.e., $q \in \text{dom } \partial h =: \mathcal{Q}$, as claimed.

(2) To establish (B.2), it suffices to show that it holds for all $q \in \mathcal{Q}$ (by continuity). To do so, let

$$\phi(t) = h(q + t(p - q)) - [h(q) + \langle \psi, q + t(p - q) \rangle]. \quad (\text{B.5})$$

Since h is strongly convex relative $\psi \in \partial h(q)$ by (B.1), it follows that $\phi(t) \geq 0$ with equality if and only if $t = 0$. Moreover, note that $\psi(t) = \langle \nabla h(q + t(p - q)) - \psi, p - q \rangle$ is a continuous selection of subgradients of ϕ . Given that ϕ and ψ are both continuous on $[0, 1]$, it follows that ϕ is continuously differentiable and $\phi' = \psi$ on $[0, 1]$. Thus, with ϕ convex and $\phi(t) \geq 0 = \phi(0)$ for all $t \in [0, 1]$, we conclude that $\phi'(0) = \langle \nabla h(q) - \psi, p - q \rangle \geq 0$, from which our claim follows.

Finally, the Fréchet differentiability of h^* is a straightforward application of the envelope theorem, which is sometimes referred to in the literature as Danskin's theorem, cf. Berge [10, Chap. 4] \square

As we mentioned in the main text, much of our analysis revolves around the energy function (13) defined by means of the Fenchel-Young inequality. To formalize this, it will be convenient to introduce a more general pairing between $p \in \mathcal{V}$ and $\psi \in \mathcal{V}^*$, known as the *Fenchel coupling*. Following [40], this is defined as

$$F(p, \psi) = h(p) + h^*(\psi) - \langle \psi, p \rangle \quad \text{for all } p \in \text{dom } h, \psi \in \mathcal{V}^*. \quad (\text{B.6})$$

The following series of lemmas gathers some basic properties of the Fenchel coupling. The first is a lower bound for the Fenchel coupling in terms of the ambient norm in \mathcal{V} :

Lemma B.2. *Let h be a regularizer on \mathcal{X} with strong convexity modulus K . Then, for all $p \in \text{dom } h$ and all $\psi \in \mathcal{V}^*$, we have*

$$F(p, \psi) \geq \frac{K}{2} \|Q(\psi) - p\|^2. \quad (\text{B.7})$$

Proof. By the definition of F and the inequality (B.2), we have:

$$\begin{aligned} F(p, \psi) &= h(p) + h^*(\psi) - \langle \psi, p \rangle = h(p) + \langle \psi, Q(\psi) \rangle - h(Q(\psi)) - \langle \psi, p \rangle \\ &\geq h(p) - h(Q(\psi)) - \langle \nabla h(\psi), Q(\psi) - p \rangle \\ &\geq \frac{K}{2} \|Q(\psi) - p\|^2 \end{aligned} \quad (\text{B.8})$$

where we used (B.2) in the second line, and the strong convexity of h in the last. \square

Our next result is the primal-dual analogue of the so-called “three-point identity” for the Bregman divergence [22]:

Proposition B.1. *Let h be a regularizer on \mathcal{X} , fix some $p \in \mathcal{V}$, $\psi, \psi^+ \in \mathcal{V}^*$, and let $q = Q(\psi)$. Then:*

$$F(p, \psi^+) = F(p, \psi) + F(q, \psi^+) + \langle \psi^+ - \psi, q - p \rangle. \quad (\text{B.9})$$

Proof. By definition:

$$\begin{aligned} F(p, \psi^+) &= h(p) + h^*(\psi^+) - \langle \psi^+, p \rangle \\ F(p, \psi) &= h(p) + h^*(\psi) - \langle \psi, p \rangle. \end{aligned} \quad (\text{B.10})$$

Thus, by subtracting the above, we get:

$$\begin{aligned} F(p, \psi^+) - F(p, \psi) &= h(p) + h^*(\psi^+) - \langle \psi^+, p \rangle - h(p) - h^*(\psi) + \langle \psi, p \rangle \\ &= h^*(\psi^+) - h^*(\psi) - \langle \psi^+ - \psi, p \rangle \\ &= h^*(\psi^+) - \langle \psi, Q(\psi) \rangle + h(Q(\psi)) - \langle \psi^+ - \psi, p \rangle \\ &= h^*(\psi^+) - \langle \psi, q \rangle + h(q) - \langle \psi^+ - \psi, p \rangle \\ &= h^*(\psi^+) + \langle \psi^+ - \psi, q \rangle - \langle \psi^+, q \rangle + h(q) - \langle \psi^+ - \psi, p \rangle \\ &= F(q, \psi^+) + \langle \psi^+ - \psi, q - p \rangle \end{aligned} \quad (\text{B.11})$$

and our proof is complete. \square

We are now in a position to state and prove a key inequality for the Fenchel coupling:

Proposition B.2. *Let h be a regularizer on \mathcal{X} with convexity modulus K , fix some $p \in \text{dom } h$, and let $q = Q(\psi)$ for some $\psi \in \mathcal{V}^*$. Then, for all $v \in \mathcal{V}^*$, we have:*

$$F(p, \psi + v) \leq F(p, \psi) + \langle v, q - p \rangle + \frac{1}{2K} \|v\|_*^2 \quad (\text{B.12})$$

Proof. Let $q = Q(\psi)$, $\psi^+ = \psi + v$, and $q^+ = Q(\psi^+)$. Then, by the three-point identity (B.9), we have

$$F(p, \psi) = F(p, \psi^+) + F(q^+, \psi) + \langle \psi - \psi^+, q^+ - p \rangle. \quad (\text{B.13})$$

Hence, after rearranging:

$$\begin{aligned} F(p, \psi^+) &= F(p, \psi) - F(q^+, \psi) + \langle v, q^+ - p \rangle \\ &= F(p, \psi) - F(q^+, \psi) + \langle v, q - p \rangle + \langle v, q^+ - q \rangle. \end{aligned} \quad (\text{B.14})$$

By Young's inequality, we also have

$$\langle v, q^+ - q \rangle \leq \frac{K}{2} \|q^+ - q\|^2 + \frac{1}{2K} \|v\|_*^2. \quad (\text{B.15})$$

Thus, substituting in (B.14), we get

$$F(p, \psi^+) \leq F(p, \psi) + \langle v, q - p \rangle + \frac{1}{2K} \|v\|_*^2 - F(q^+, \psi) + \frac{K}{2} \|q^+ - q\|^2. \quad (\text{B.16})$$

Our claim then follows by noting that $F(q^+, \psi) \geq \frac{K}{2} \|q^+ - q\|^2$ (cf. Lemma B.2 above). \square

APPENDIX C. REGRET DERIVATIONS

Notation: from losses to payoffs. In this appendix, we prove the general regret guarantees for Alg. 1. For notational convenience, we will switch in what follows from “losses” to “payoffs”, i.e., we will assume that the learner is encountering a sequence of payoff functions $u_t = -\ell_t$ and gets as feedback the model $\hat{u}_t = -\hat{u}_t$.

C.1. Basic bounds and preliminaries. We begin by providing some template regret bounds that we will use as a toolkit in the sequel. As a warm-up, we prove the basic comparison lemma between simple and pure strategies:

Lemma 1. *Let \mathcal{U} be a convex neighborhood of x in \mathcal{K} and let $\chi \in \mathcal{X}$ be a simple strategy supported on \mathcal{U} . Then, $\text{Reg}_x(T) \leq \text{Reg}_\chi(T) + L \text{diam}(\mathcal{U})T$.*

Proof. By Assumption 1, we have $u_t(x) \leq u_t(x') + L\|x - x'\| \leq u_t(x') + L \text{diam}(\mathcal{U})$ for all $x' \in \mathcal{U}$. Hence, taking expectations on both sides relative to χ , we get $u_t(x) \leq \langle u_t, \chi \rangle + L \text{diam}(\mathcal{U})$. Our claim then follows by summing over $t = 1, 2, \dots, T$ and invoking the definition of the regret. \square

We now turn to the derivation of our main regret guarantees as outlined in Section 4. Much of the analysis to follow will revolve around the energy function (13) which, for convenience, we restate below in terms of the Fenchel coupling (B.6):

$$E_t := \frac{1}{\eta_t} [h(\chi) + h^*(\eta_t y_t) - \langle \eta_t y_t, \chi \rangle] = \frac{1}{\eta_t} F(\chi, \eta_t y_t). \quad (\text{13})$$

In words, E_t essentially measures the primal-dual “distance” between the benchmark strategy χ and the aggregate model y_t , taking into account the inflation of the latter by η_t in (DA).

Our overall proof strategy will then be to relate the regret incurred by the optimizer to the evolution of E_t over time. To that end, an application of Abel's summation formula gives:

$$\begin{aligned} E_{t+1} - E_t &= \frac{1}{\eta_{t+1}} F(\chi, \eta_{t+1} y_{t+1}) - \frac{1}{\eta_t} F(\chi, \eta_t y_t) \\ &= \frac{1}{\eta_{t+1}} F(\chi, \eta_{t+1} y_{t+1}) - \frac{1}{\eta_t} F(\chi, \eta_t y_{t+1}) \end{aligned} \quad (\text{C.1a})$$

$$+ \frac{1}{\eta_t} F(\chi, \eta_t y_{t+1}) - \frac{1}{\eta_t} F(\chi, \eta_t y_t). \quad (\text{C.1b})$$

We now proceed to unpack the two terms (C.1a) and (C.1b) separately, beginning with the latter.

To do so, substituting $p \leftarrow \chi$, $\psi \leftarrow \eta_t y_t$ and $\psi^+ \leftarrow \eta_t y_{t+1}$ in Proposition B.1 yields

$$\begin{aligned} (\text{C.1b}) &= \frac{1}{\eta_t} [F(\chi, \eta_t y_t + \eta_t u_t) - F(\chi, \eta_t y_t)] \\ &= \frac{1}{\eta_t} [F(p_t, \eta_t y_{t+1}) + \langle \eta_t u_t, p_t - \chi \rangle] \\ &= \frac{F(p_t, \eta_t y_{t+1})}{\eta_t} + \langle u_t, p_t - \chi \rangle \end{aligned} \quad (\text{C.2})$$

where we used the definition $p_t = Q(\eta_t y_t)$ of p_t . We thus obtain the interim expression

$$E_{t+1} = E_t + (\text{C.1a}) + \langle u_t, p_t - \chi \rangle + \frac{F(p_t, \eta_t y_{t+1})}{\eta_t} \quad (\text{C.3})$$

Moving forward, for the term (C.1a), the definition of the Fenchel coupling (B.6) readily yields:

$$(\text{C.1a}) = \left[\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right] h(\chi) + \frac{1}{\eta_{t+1}} h^*(\eta_{t+1} y_{t+1}) - \frac{1}{\eta_t} h^*(\eta_t y_{t+1}). \quad (\text{C.4})$$

Consider now the function $\varphi(\eta) = \eta^{-1} [h^*(\eta\psi) + \min h]$ for arbitrary $\psi \in \mathcal{L}^\infty(\mathcal{K})$. By Lemma B.1, h^* is Fréchet differentiable with $D_v h^*(\cdot) = \langle v, Q(\cdot) \rangle$ for all $v \in \mathcal{V}^*$, so a simple differentiation yields

$$\begin{aligned} \varphi'(\eta) &= \frac{1}{\eta} \langle \psi, Q(\eta\psi) \rangle - \frac{1}{\eta^2} [h^*(\eta\psi) + \min h] \\ &= \frac{1}{\eta^2} [\langle \eta\psi, Q(\eta\psi) \rangle - h^*(\eta\psi) - \min h] \\ &= \frac{1}{\eta^2} [h(Q(\eta\psi)) - \min h] \geq 0, \end{aligned} \quad (\text{C.5})$$

where we used the Fenchel-Young inequality as an equality in the second-to-last line. Since $\eta_{t+1} \leq \eta_t$, the above shows that $\varphi(\eta_t) \geq \varphi(\eta_{t+1})$. Hence, substituting $\psi \leftarrow y_{t+1}$, we ultimately obtain

$$\frac{1}{\eta_{t+1}} h^*(\eta_{t+1} y_{t+1}) - \frac{1}{\eta_t} h^*(\eta_t y_{t+1}) \leq \left[\frac{1}{\eta_t} - \frac{1}{\eta_{t+1}} \right] \min h. \quad (\text{C.6})$$

Therefore, combining (C.3) and (C.6), we have proved the following template bound:

Lemma C.1. *For all $\chi \in \mathcal{X}$, the policy (DA) enjoys the bound*

$$E_{t+1} \leq E_t + \langle u_t, p_t - \chi \rangle + \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) [h(\chi) - \min h] + \frac{1}{\eta_t} F(p_t, \eta_t y_{t+1}). \quad (\text{C.7})$$

We are now in a position to prove our basic energy inequality (restated below for convenience):

Lemma 2. For all $\chi \in \mathcal{X}$, we have:

$$E_{t+1} \leq E_t + \langle u_t, \chi - p_t \rangle + (\eta_{t+1}^{-1} - \eta_t^{-1})[h(\chi) - \min h] + \frac{\eta_t}{2K} \|u_t\|_*^2. \quad (14)$$

Proof. Going back to [Proposition B.2](#) and setting $p \leftarrow p_t$, $\psi \leftarrow \eta_t y_t$ and $v \leftarrow \eta_t u_t$, we get

$$F(p_t, \eta_t y_{t+1}) \leq F(p_t, \eta_t y_t) + \langle \eta_t u_t, p_t - p_t \rangle + \frac{\eta_t^2}{2K} \|u_t\|_*^2 = \frac{\eta_t^2}{2K} \|u_t\|_*^2 \quad (C.8)$$

where we used the fact that $p_t = Q(\eta_t y_t)$. Our claim then follows by dividing both sides by η_t and substituting in [Lemma C.1](#). \square

We will come back to these results as needed.

C.2. Static regret guarantees. We are now ready to prove our static regret results for [Alg. 1](#). We begin with the precursor to our main result in that respect:

Proposition 1. For any simple strategy $\chi \in \mathcal{X}$, [Alg. 1](#) enjoys the bound

$$\text{Reg}_\chi(T) \leq \eta_{T+1}^{-1} [h(\chi) - \min h] + \sum_{t=1}^T \langle e_t, \chi - p_t \rangle + \frac{1}{2K} \sum_{t=1}^T \eta_t \|u_t\|_*^2. \quad (12)$$

Proof. Recalling the decomposition $u_t = u_t + e_t$ for the learner's inexact models, a simple rearrangement of [Lemma 2](#) gives

$$\langle u_t, \chi - p_t \rangle \leq E_t - E_{t+1} + \langle e_t, p_t - \chi \rangle + (\eta_{t+1}^{-1} - \eta_t^{-1})[h(\chi) - \min h] + \frac{\eta_t}{2K} \|u_t\|_*^2. \quad (C.9)$$

Thus, telescoping over $t = 1, 2, \dots, T$, we get

$$\begin{aligned} \text{Reg}_\chi(T) &\leq E_1 - E_{T+1} + \left(\frac{1}{\eta_{T+1}} - \frac{1}{\eta_1} \right) [h(\chi) - \min h] + \sum_{t=1}^T \langle e_t, p_t - \chi \rangle + \frac{1}{2K} \sum_{t=1}^T \eta_t \|u_t\|_*^2 \\ &\leq \frac{h(\chi) - \min h}{\eta_{T+1}} + \sum_{t=1}^T \langle e_t, p_t - \chi \rangle + \frac{1}{2K} \sum_{t=1}^T \eta_t \|u_t\|_*^2, \end{aligned} \quad (C.10)$$

where we used the fact that $E_t \geq 0$ for all t and $E_1 = \eta_1^{-1} [h(\chi) + h^*(0)] = \eta_1^{-1} [h(\chi) - \min h]$. \square

As a simple application of [Lemma 2](#), we get the following bound for simple comparators:

Corollary 3. For all $\chi \in \mathcal{X}$, [Alg. 1](#) guarantees

$$\mathbb{E}[\text{Reg}_\chi(T)] \leq \frac{h(\chi) - \min h}{\eta_{T+1}} + 2 \sum_{t=1}^T B_t + \frac{1}{2K} \sum_{t=1}^T \eta_t \mathbb{E}[\|u_t\|_*^2], \quad (C.11)$$

Proof. Simply take expectations over [\(12\)](#) and use the fact that

$$\mathbb{E}[\langle e_t, p_t - \chi \rangle] = \mathbb{E}[\langle \mathbb{E}[e_t | \mathcal{F}_t], p_t - \chi \rangle] = \mathbb{E}[\langle b_t, p_t - \chi \rangle] \leq \mathbb{E}[\|b_t\|_\infty \|p_t - \chi\|_1] \leq 2B_t. \quad \square$$

We are finally in a position to prove the main static regret guarantee of [Alg. 1](#):

Theorem 1. Fix $x \in \mathcal{K}$, let \mathcal{C} be a convex neighborhood of x in \mathcal{K} , and suppose that [Alg. 1](#) is run with a decomposable regularizer $h(p) = \int_{\mathcal{K}} \theta \circ p$. Then, letting $\phi(z) = z\theta(1/z)$ for $z > 0$, we have:

$$\mathbb{E}[\text{Reg}_x(T)] \leq \frac{\phi(\lambda(\mathcal{C})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + L \text{diam}(\mathcal{C})T + 2 \sum_{t=1}^T B_t + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t M_t^2. \quad (15)$$

In particular, if [Alg. 1](#) is run with learning rate $\eta_t \propto 1/t^\rho$, $\rho \in (0, 1)$, and inexact models such that $B_t = \mathcal{O}(1/t^\beta)$ and $M_t^2 = \mathcal{O}(t^{2\mu})$ for some $\beta, \mu \geq 0$, we have:

$$\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(\phi(T^{-n\kappa})T^\rho + T^{1-\kappa} + T^{1-\beta} + T^{1+2\mu-\rho}) \quad \text{for all } \kappa \geq 0. \quad (16)$$

Proof. To simplify the proof, we will make the normalizing assumption $\theta(0) = 0$; if this is not the case, θ can always be shifted by $\theta(0)$ for this condition to hold. [Note that [Examples 4](#) and [5](#) both satisfy this convention.]

With this in mind, let \mathcal{C} be a convex neighborhood of x in \mathcal{K} , and let $\text{unif}_{\mathcal{C}} = \lambda(\mathcal{C})^{-1} \mathbb{1}_{\mathcal{C}}$ denote the (simple) strategy that assigns uniform probability to the elements of \mathcal{C} and zero to all other points in \mathcal{C} . We then have:

$$h(\text{unif}_{\mathcal{C}}) = \int_{\mathcal{K}} \theta(\text{unif}_{\mathcal{C}}) = \int_{\mathcal{K}} \theta(\mathbb{1}_{\mathcal{C}} / \lambda(\mathcal{C})) = \int_{\mathcal{C}} \theta(1/\lambda(\mathcal{C})) = \lambda(\mathcal{C})\theta(1/\lambda(\mathcal{C})) = \phi(\lambda(\mathcal{C})). \quad (\text{C.12})$$

Moreover, since h is decomposable and the probability constraint $\int_{\mathcal{K}} \chi = 1$ is symmetric, the minimum of h over \mathcal{X} will be attained at the uniform strategy $\text{unif}_{\mathcal{K}} = \lambda(\mathcal{K})^{-1} \mathbb{1}_{\mathcal{K}}$. Thus, with \mathcal{X} weakly dense in $\text{dom } h$, we obtain

$$\min h = h(\text{unif}_{\mathcal{K}}) = \int_{\mathcal{K}} \theta(\mathbb{1}_{\mathcal{K}} / \lambda(\mathcal{K})) = \phi(\lambda(\mathcal{K})). \quad (\text{C.13})$$

In view of all this, [Corollary 3](#) applied to $\chi = \text{unif}_{\mathcal{C}}$ yields

$$\mathbb{E}[\text{Reg}_{\chi}(T)] \leq \frac{\phi(\lambda(\mathcal{C})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + 2 \sum_{t=1}^T B_t + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t M_t^2, \quad (\text{C.14})$$

where we used the fact that $\|\cdot\|_{\text{TV}} \leq \alpha \|\cdot\|$ so $\|\cdot\|_* \leq \alpha \|\cdot\|_{\infty}$. The bound [\(15\)](#) then follows by combining the above with [Lemma 1](#).

Regarding the bound [\(16\)](#), we first note that this is not a pseudo-regret bound but a bona fide bound for the learner's *expected regret* (so we cannot simply our point-dependent bound over $x \in \mathcal{K}$). In light of this, our first step will be to consider a ‘‘uniform’’ simple approximant for every $x \in \mathcal{K}$. To that end, building on an idea by Blum & Kalai [\[12\]](#) and Krichene et al. [\[38\]](#), fix a shrinkage factor $\delta > 0$ and let $\mathcal{K}_{\delta}(x) = \{x + \delta(x' - x) : x' \in \mathcal{K}\} \subseteq \mathcal{K}$ denote the homothetic transformation that shrinks \mathcal{K} to a fraction δ of its original size and then transports it to $x \in \mathcal{K}$. By construction, we have $x \in \mathcal{K}_{\delta}(x) \subseteq \mathcal{K}$ and, moreover, $\text{diam}(\mathcal{K}_{\delta}(x)) = \delta \text{diam}(\mathcal{K})$ and $\lambda(\mathcal{K}_{\delta}(x)) = \delta^n \lambda(\mathcal{K})$. Then, letting $\mu_x := \text{unif}_{\mathcal{K}_{\delta}(x)}$ denote the uniform strategy supported on $\mathcal{K}_{\delta}(x)$, we get

$$\mathbb{E}[\text{Reg}(T)] = \mathbb{E} \left[\max_{x \in \mathcal{K}} \text{Reg}_x(T) \right] \leq \mathbb{E} \left[\max_{x \in \mathcal{K}} \text{Reg}_{\mu_x}(T) \right] + \delta L \text{diam}(\mathcal{K})T, \quad (\text{C.15})$$

where, in the last step, we used [Lemma 1](#).

Now, by [Proposition 1](#), we have

$$\begin{aligned} \text{Reg}_{\mu_x}(T) &\leq \frac{h(\mu_x) - \min h}{\eta_{T+1}} + \sum_{t=1}^T \langle e_t, p_t - \mu_x \rangle + \frac{1}{2K} \sum_{t=1}^T \eta_t \|u_t\|_*^2 \\ &\leq \frac{\phi(\delta^n \lambda(\mathcal{K})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + \sum_{t=1}^T \langle e_t, p_t - \mu_x \rangle + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t \|u_t\|_{\infty}^2. \end{aligned} \quad (\text{C.16})$$

and hence

$$\mathbb{E} \left[\max_{x \in \mathcal{K}} \text{Reg}_{\mu_x}(T) \right] \leq \frac{\phi(\delta^n \lambda(\mathcal{K})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + \mathbb{E} \left[\max_{x \in \mathcal{K}} \sum_{t=1}^T \langle e_t, p_t - \mu_x \rangle \right] + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t M_t^2. \quad (\text{C.17})$$

Thus, to proceed, it suffices to bound the second term of the above expression.

To do so, introduce the auxiliary process

$$\tilde{y}_{t+1} = \tilde{y}_t - z_t, \quad \tilde{p}_{t+1} = Q(\eta_{t+1} \tilde{y}_{t+1}), \quad (\text{C.18})$$

with $\tilde{p}_1 = p_1$. We then have

$$\begin{aligned} \sum_{t=1}^T \langle e_t, p_t - \mu_x \rangle &= \sum_{t=1}^T \langle e_t, (p_t - \tilde{p}_t) + (\tilde{p}_t - \mu_x) \rangle \\ &= \sum_{t=1}^T \langle e_t, p_t - \tilde{p}_t \rangle + \sum_{t=1}^T \langle b_t, \tilde{p}_t - \mu_x \rangle + \sum_{t=1}^T \langle z_t, \tilde{p}_t - \mu_x \rangle \end{aligned} \quad (\text{C.19})$$

so it suffices to derive a bound for each of these terms. This can be done as follows:

- (1) The first term of (C.19) does not depend on x , so we have

$$\mathbb{E} \left[\max_{x \in \mathcal{K}} \sum_{t=1}^T \langle e_t, p_t - \tilde{p}_t \rangle \right] = \sum_{t=1}^T \mathbb{E} [\mathbb{E}[\langle e_t, p_t - \tilde{p}_t \rangle \mid \mathcal{F}_t]] \leq 2B_t \quad (\text{C.20})$$

where, in the last step, we used the definition (7a) of B_t and the bound

$$\langle b_t, p_t - \tilde{p}_t \rangle \leq \|p_t - \tilde{p}_t\|_1 \|b_t\|_\infty \leq 2B_t. \quad (\text{C.21})$$

- (2) The second term of (C.19) can be similarly bounded as

$$\mathbb{E} \left[\max_{x \in \mathcal{K}} \sum_{t=1}^T \langle b_t, \tilde{p}_t - \mu_x \rangle \right] \leq \mathbb{E} [\|\tilde{p}_t - \mu_x\|_1 \|b_t\|_*] \leq 2B_t. \quad (\text{C.22})$$

- (3) The third term is more challenging; the main idea will be to apply [Proposition 1](#) on the sequence $-z_t$, $t = 1, 2, \dots$, viewed itself as a sequence of virtual payoff functions. Doing just that, we get:

$$\begin{aligned} \sum_{t=1}^T \langle z_t, \tilde{p}_t - \mu_x \rangle &\leq \frac{h(\mu_x) - \min h}{\eta_{T+1}} + \frac{1}{2K} \sum_{t=1}^T \eta_t \|z_t\|_*^2 \\ &\leq \frac{\phi(\delta^n \lambda(\mathcal{K})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t \|z_t\|_\infty^2. \end{aligned} \quad (\text{C.23})$$

Thus, after maximizing and taking expectations, we obtain

$$\mathbb{E} \left[\max_{x \in \mathcal{K}} \langle z_t, \tilde{p}_t - \mu_x \rangle \right] \leq \frac{\phi(\delta^n \lambda(\mathcal{K})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t \sigma_t^2. \quad (\text{C.24})$$

Therefore, plugging [Eqs. \(C.20\), \(C.22\) and \(C.24\)](#) into [\(C.19\)](#) and substituting the result to [\(C.17\)](#), we finally get

$$\mathbb{E} \left[\max_{x \in \mathcal{K}} \text{Reg}_{\mu_x}(T) \right] \leq 2 \frac{\phi(\delta^n \lambda(\mathcal{K})) - \phi(\lambda(\mathcal{K}))}{\eta_{T+1}} + \frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t (M_t^2 + \sigma_t^2). \quad (\text{C.25})$$

The guarantee [\(16\)](#) then follows by taking $\delta^n \lambda(\mathcal{K}) = T^{-n\kappa}$ for some $\kappa \geq 0$ and plugging everything back in [\(C.15\)](#). \square

C.3. Dynamic regret guarantees. We now turn to the algorithm's dynamic regret guarantees, as encoded by [Theorem 2](#) (stated below for convenience):

Theorem 2. *Suppose that the Hedge variant of [Alg. 1](#) is run with learning rate $\eta_t \propto 1/t^\rho$ and inexact models with $B_t = \mathcal{O}(1/t^\beta)$ and $M_t^2 = \mathcal{O}(t^{2\mu})$ for some $\beta, \mu \geq 0$. Then:*

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1+2\mu-\rho} + T^{1-\beta} + T^{2\rho-2\mu} V_T). \quad (\text{19})$$

In particular, if $V_T = \mathcal{O}(T^\nu)$ for some $\nu < 1$ and the learner's feedback is unbiased and bounded in mean square (i.e., $B_t = 0$ and $\sup_t M_t < \infty$), the choice $\rho = (1 - \nu)/3$ guarantees

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{\frac{2+\nu}{3}}). \quad (20)$$

Proof of Theorem 2. As we discussed in the main body of our paper, our proof strategy will be to decompose the horizon of play into m virtual segments, estimate the learner's regret over each segment, and then compare the learner's regret *per-segment* to the corresponding dynamic regret over said segment. We stress here again that this partition is only made for the sake of the analysis, and does not involve restarting the algorithm – e.g., as in Besbes et al. [11].

To make this precise, we first partition the interval $\mathcal{T} = [1..T]$ into m contiguous segments \mathcal{T}_k , $k = 1, \dots, m$, each of length Δ (except possibly the m -th one, which might be smaller). More explicitly, take the window length to be of the form $\Delta = \lceil T^\gamma \rceil$ for some constant $\gamma \in [0, 1]$ to be determined later. In this way, the number of windows is $m = \lceil T/\Delta \rceil = \Theta(T^{1-\gamma})$ and the k -th window will be of the form $\mathcal{T}_k = [(k-1)\Delta + 1..k\Delta]$ for all $k = 1, \dots, m-1$ (the value $k = m$ is excluded as the m -th window might be smaller). For concision, we will denote the learner's static regret over the k -th window as $\text{Reg}(\mathcal{T}_k) = \max_{x \in \mathcal{K}} \sum_{t \in \mathcal{T}_k} \langle u_t, \delta_x - p_t \rangle$ (and likewise for its dynamic counterpart).

To proceed, let $\mathcal{S} \subseteq \mathcal{T}$ be a sub-interval of \mathcal{T} and write $x_{\mathcal{S}}^* \in \arg \max_{x \in \mathcal{K}} \sum_{s \in \mathcal{S}} u_s(x)$ for any action that is optimal on average over the interval \mathcal{S} . To ease notation, we also write $x_t^* \equiv x_{\{t\}}^* \in \arg \max_{x \in \mathcal{K}} u_t(x)$ for any action that is optimal at time t , and $x_k^* \equiv x_{\mathcal{T}_k}^*$ for any action that is optimal on average over the k -th window. Then, for all $t \in \Delta_k$, $k = 1, 2, \dots, m$, we have

$$\langle u_t, \delta_{x_t^*} - p_t \rangle = \langle u_t, \delta_{x_k^*} - p_t \rangle + [u_t(x_t^*) - u_t(x_k^*)] \quad (C.26)$$

so the learner's dynamic regret over \mathcal{T}_k can be bounded as

$$\text{DynReg}(\mathcal{T}_k) = \sum_{t \in \mathcal{T}_k} \langle u_t, \delta_{x_t^*} - p_t \rangle + \sum_{t \in \mathcal{T}_k} [u_t(x_t^*) - u_t(x_k^*)] = \text{Reg}(\mathcal{T}_k) + \sum_{t \in \mathcal{T}_k} [u_t(x_t^*) - u_t(x_k^*)]. \quad (C.27)$$

Following a batch-comparison technique originally due to Besbes et al. [11], let $\tau_k = \min \mathcal{T}_k$ denote the beginning of the k -th window, and let $x_{\tau_k}^*$ denote a maximizer of the first payoff function encountered in the window \mathcal{T}_k (this choice could of course be arbitrary). Thus, given that x_k^* maximizes the per-window aggregate $\sum_{t \in \mathcal{T}_k} u_t(x)$, we obtain:

$$\begin{aligned} \sum_{t \in \mathcal{T}_k} [u_t(x_t^*) - u_t(x_k^*)] &\leq \sum_{t \in \mathcal{T}_k} [u_t(x_t^*) - u_t(x_{\tau_k}^*)] \\ &\leq |\mathcal{T}_k| \max_{t \in \mathcal{T}_k} [u_t(x_t^*) - u_t(x_{\tau_k}^*)] \leq 2\Delta V_k, \end{aligned} \quad (C.28)$$

where we let $V_k = \sum_{t \in \mathcal{T}_k} \|u_{t+1} - u_t\|_\infty$. In turn, combining (C.28) with (C.27), we get:

$$\text{DynReg}(\mathcal{T}_k) \leq \text{Reg}(\mathcal{T}_k) + 2\Delta V_k, \quad (C.29)$$

and hence, after summing over all windows:

$$\text{DynReg}(T) \leq \sum_{k=1}^m \text{Reg}(\mathcal{T}_k) + 2\Delta V_T. \quad (C.30)$$

Now Theorem 1 applied to the Hedge variant of Alg. 1 readily yields

$$\mathbb{E}[\text{Reg}(\mathcal{T}_k)] = \mathcal{O} \left((k\Delta)^\rho + \Delta^{1-\kappa} + \sum_{t \in \mathcal{T}_k} t^{-\beta} + \sum_{t \in \mathcal{T}_k} t^{1+2\mu-\rho} \right) \quad (C.31)$$

so, after summing over all windows, we have

$$\begin{aligned} \sum_{k=1}^m \mathbb{E}[\text{Reg}(\mathcal{T}_k)] &= \mathcal{O}\left(\Delta^\rho \sum_{k=1}^m k^\rho + m\Delta^{1-\kappa} + \sum_{t=1}^T t^{-\beta} + \sum_{t=1}^T t^{2\mu-\rho}\right) \\ &= \mathcal{O}(\Delta^\rho m^{1+\rho} + m\Delta^{1-\kappa} + T^{1-\beta} + T^{1+2\mu-\rho}). \end{aligned} \quad (\text{C.32})$$

Since $\Delta = \mathcal{O}(T^\gamma)$ and $m = \mathcal{O}(T/\Delta) = \mathcal{O}(T^{1-\gamma})$, we get

$$\Delta^\rho m^{1+\rho} = \mathcal{O}((m\Delta)^\rho m) = \mathcal{O}(T^{\gamma\rho} T^{(1-\gamma)(1+\rho)}) = \mathcal{O}(T^{1+\rho-\gamma}) \quad (\text{C.33})$$

and, likewise

$$m\Delta^{1-\kappa} = \mathcal{O}(T\Delta^{-\kappa}) = \mathcal{O}(TT^{-\gamma\kappa}) = \mathcal{O}(T^{1-\gamma\kappa}). \quad (\text{C.34})$$

Then, substituting in (C.32) and (C.30), we finally get the dynamic regret bound

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1+\rho-\gamma} + T^{1-\gamma\kappa} + T^{1-\beta} + T^{1+2\mu-\rho} + T^\gamma V_T). \quad (\text{C.35})$$

To balance the above expression, we take $\gamma = 2\rho - 2\mu$ for the window size exponent (which calibrates the first and fourth terms in the sum above) and $\kappa = \beta/\gamma = \beta/(2\rho - 2\mu)$ (for the second and the third). In this way, we finally obtain

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1-\beta} + T^{1+2\mu-\rho} + T^{2\rho-2\mu} V_T) \quad (\text{C.36})$$

and our proof is complete. \square

APPENDIX D. DERIVATIONS FOR THE BANDIT FRAMEWORK

In this appendix, we aim at deriving guarantees for the Hedge variant of Alg. 2 using template bounds from Appendix C. We start by stating preliminary results that are used in the sequel.

D.1. Preliminary results. We first present a technical bound for the convex conjugate of the entropic regularizer (more on this below):

Lemma D.1. *For all $\psi, v \in \mathcal{V}^*$, there exists $\xi \in [0, 1]$ such that:*

$$\log\left(\int_{\mathcal{K}} \exp(\psi + v)\right) \leq \log\left(\int_{\mathcal{K}} \exp(\psi)\right) + \langle v, \Lambda(\psi) \rangle + \frac{1}{2} \langle v^2, \Lambda(\psi + \xi v) \rangle. \quad (\text{D.1})$$

Proof. Consider the function $\phi: [0, 1] \rightarrow \mathbb{R}$ with $\phi(t) = \log\left(\int_{\mathcal{K}} \exp(\psi + tv)\right)$. By construction, $\phi(0) = \log\left(\int_{\mathcal{K}} \exp(\psi)\right)$ and $\phi(1) = \log\left(\int_{\mathcal{K}} \exp(\psi + v)\right)$. Thus, by a second-order Taylor expansion with Lagrange remainder, we have:

$$\phi(1) = \phi(0) + \phi'(0) + \frac{1}{2} \phi''(\xi) \quad (\text{D.2})$$

for some $\xi \in [0, 1]$.

Now, for all $t \in [0, 1]$, $\phi'(t) = \frac{\int_{\mathcal{K}} v \exp(\psi + tv)}{\int_{\mathcal{K}} \exp(\psi + tv)}$, which in turns gives

$$\phi'(0) = \frac{\int_{\mathcal{K}} v \exp(\psi)}{\int_{\mathcal{K}} \exp(\psi)} = \langle v, \Lambda(\psi) \rangle. \quad (\text{D.3})$$

As for the second order derivative of ϕ , we have for all $t \in [0, 1]$:

$$\begin{aligned} \phi''(t) &= \frac{\partial}{\partial t} \left[\frac{\int_{\mathcal{K}} v \exp(\psi + tv)}{\int_{\mathcal{K}} \exp(\psi + tv)} \right] \\ &= \frac{\int_{\mathcal{K}} v^2 \exp(\psi + tv) \int_{\mathcal{K}} \exp(\psi + tv) - \left(\int_{\mathcal{K}} v \exp(\psi + tv)\right)^2}{\left(\int_{\mathcal{K}} \exp(\psi + tv)\right)^2} \end{aligned}$$

$$\leq \frac{\int_{\mathcal{K}} v^2 \exp(\psi + tv) \int_{\mathcal{K}} \exp(\psi + tv)}{(\int_{\mathcal{K}} \exp(\psi + tv))^2} = \frac{\int_{\mathcal{K}} v^2 \exp(\psi + tv)}{\int_{\mathcal{K}} \exp(\psi + tv)} \quad (\text{D.4})$$

Thus, for all $t \in [0, 1]$, we get

$$\phi''(t) \leq \langle v^2, \Lambda(\psi + tv) \rangle. \quad (\text{D.5})$$

Our claim then follows by injecting (D.3) and (D.5) into (D.2). \square

In the next lemma, we now present an expression of the Fenchel coupling in the specific case of the negentropy regularizer $h(p) = \int_{\mathcal{K}} p \log p$.

Lemma D.2. *In the case of the negentropy regularizer $h(p) = \int_{\mathcal{K}} p \log p$, the Fenchel coupling for all $\psi \in \mathcal{V}^*$ and $p \in \text{dom } h$ is given by*

$$F(p, \psi) = \int_{\mathcal{K}} p \log p + \log \left(\int_{\mathcal{K}} \exp(\psi) \right) - \langle \psi, p \rangle. \quad (\text{D.6})$$

Proof. We remind the general expression of the Fenchel coupling given in (B.6):

$$F(p, \psi) = h(p) + h^*(\psi) - \langle \psi, p \rangle \quad \text{for all } p \in \text{dom } h, \psi \in \mathcal{V}^*, \quad (\text{D.7})$$

where $h^*(\psi) = \max_{p \in \mathcal{V}} \{\langle \psi, p \rangle - h(p)\}$. In the case of the negentropy regularizer $h(p) = \int_{\mathcal{K}} p \log p$, we have that $\arg \max_{p \in \mathcal{V}} \{\langle \psi, p \rangle - h(p)\} = \Lambda(\psi)$ and

$$h^*(\psi) = \langle \psi, \Lambda(\psi) \rangle - h(\Lambda(\psi)). \quad (\text{D.8})$$

Combining the above, we then get:

$$\begin{aligned} h^*(\psi) &= \int_{\mathcal{K}} \psi \Lambda(\psi) - \int_{\mathcal{K}} \Lambda(\psi) \log \Lambda(\psi) \\ &= \int_{\mathcal{K}} \psi \Lambda(\psi) - \int_{\mathcal{K}} \Lambda(\psi) \psi + \int_{\mathcal{K}} \log \left(\int_{\mathcal{K}} \exp(\psi) \right) \Lambda(\psi) \\ &= \log \left(\int_{\mathcal{K}} \exp(\psi) \right). \end{aligned}$$

which, combined with (B.6), delivers (D.6). \square

Finally we state a result enabling to control the difference between the regret $\text{Reg}(T)$ and $\widetilde{\text{Reg}}(T)$ induced respectively by two policies p_t and \tilde{p}_t against the same rewards and models.

Lemma D.3. *For $t = 1, \dots, T$, let p_t, \tilde{p}_t be two policies with respective regret $\text{Reg}(T)$ and $\widetilde{\text{Reg}}(T)$ against a given sequence of models $(u_t)_t$ for the rewards $(u_t)_t$. Then:*

$$\text{Reg}(T) \leq \widetilde{\text{Reg}}(T) + \sum_{t=1}^T \|p_t - \tilde{p}_t\|_{\infty}. \quad (\text{D.9})$$

Proof. See Slivkins [47, Chap. 6]. \square

D.2. Hedge-specific bounds. We are now ready to adapt the template bound of Lemma C.1 to the Hedge case.

Lemma D.4. *Assuming the regularizer h is the negentropy $h(p) = \int_{\mathcal{K}} p \log p$, and that the mirror map Q corresponds to the logit operator Λ , there exists $\xi \in [0, 1]$ such that, for all $\chi \in \mathcal{X}$ the policy (DA) enjoys the bound:*

$$E_{t+1} \leq E_t + \langle u_t, p_t - \chi \rangle + \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) [h(\chi) - \min h] + \frac{\eta_t}{2} G_t(\xi)^2. \quad (\text{D.10})$$

where for all $\xi \in [0, 1]$, $G_t(\xi)^2 = \langle \Lambda(\eta_t y_t + \xi \eta_t u_t), u_t^2 \rangle$.

Proof. We know from [Lemma D.4](#) that the policy (DA) enjoys the bound:

$$E_{t+1} \leq E_t + \langle u_t, p_t - \chi \rangle + \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) [h(\chi) - \min h] + \frac{1}{\eta_t} F(p_t, \eta_t y_{t+1}). \quad (\text{D.11})$$

The following lemma will help us handle the Fenchel coupling term in [\(D.11\)](#)

Lemma D.5. *For a given t in the policy (DA), there exists $\xi \in [0, 1]$ such that the following bounds holds:*

$$F(p_t, \eta_t y_{t+1}) \leq \frac{\eta_t^2}{2} G_t(\xi)^2. \quad (\text{D.12})$$

Injecting the result given in [Lemma D.5](#) in [Eq. \(D.11\)](#) yields the stated claim. \square

Moving forward, we are only left to prove [Lemma D.5](#).

Proof. Since we are in the case of the negentropy regularizer, [Lemma D.2](#) enables to rewrite the Fenchel coupling term of [\(D.11\)](#) as:

$$F(p_t, \eta_t y_{t+1}) = \int_{\mathcal{K}} p_t \log p_t + \log \left(\int_{\mathcal{K}} \exp(\eta_t y_{t+1}) \right) - \langle \eta_t y_{t+1}, p_t \rangle. \quad (\text{D.13})$$

Injecting $y_{t+1} = y_t + u_t$ in [\(D.13\)](#) yields:

$$\begin{aligned} F(p_t, \eta_t y_{t+1}) &= \int_{\mathcal{K}} p_t \log p_t + \log \left(\int_{\mathcal{K}} \exp(\eta_t y_t + \eta_t u_t) \right) - \langle \eta_t y_t, p_t \rangle - \langle \eta_t u_t, p_t \rangle \\ &= F(p_t, \eta_t y_t) + \left(\log \int_{\mathcal{K}} \exp(\eta_t y_t + \eta_t u_t) - \log \int_{\mathcal{K}} \exp(\eta_t y_t) \right) - \langle \eta_t u_t, p_t \rangle \\ &= \log \left(\int_{\mathcal{K}} \exp(\eta_t y_t + \eta_t u_t) \right) - \log \left(\int_{\mathcal{K}} \exp(\eta_t y_t) \right) - \langle \eta_t u_t, p_t \rangle \end{aligned} \quad (\text{D.14})$$

where we used the fact that $F(p_t, \eta_t y_t) = 0$.

Now, by [Lemma D.1](#) applied to $\psi \leftarrow \eta_t y_t$ and $v \leftarrow \eta_t u_t$, there exists $\xi \in [0, 1]$ such that

$$\log \left(\int_{\mathcal{K}} \exp(\eta_t y_t + \eta_t u_t) \right) \leq \log \left(\int_{\mathcal{K}} \exp(\eta_t y_t) \right) + \eta_t \langle u_t, \Lambda(\eta_t y_t) \rangle + \frac{\eta_t^2}{2} \langle u_t^2, \Lambda(\eta_t y_t + \xi \eta_t u_t) \rangle, \quad (\text{D.15})$$

where we used the fact that $p_t = \Lambda(\eta_t y_t)$. Our claim then follows by injecting [\(D.15\)](#) into our prior expression for the Fenchel coupling $F(p_t, \eta_t y_{t+1})$ in the case of the Hedge variant. \square

Proposition D.1. *If we run the Hedge variant of [Alg. 1](#), there exists a sequence $\xi_t \in [0, 1]$ such that:*

$$\mathbb{E}[\text{Reg}_x(T)] \leq \frac{\log(\lambda(\mathcal{K})/\lambda(\mathcal{C}))}{\eta_{T+1}} + L \text{diam}(\mathcal{C})T + 2 \sum_{t=1}^T B_t + \frac{1}{2} \sum_{t=1}^T \eta_t \mathbb{E}[G_t(\xi_t)^2 | \mathcal{F}_t], \quad (\text{D.16})$$

where \mathcal{C} is a convex neighborhood of x in \mathcal{K} .

Proof. This result is obtained by using the template bound given in [Lemma D.4](#), then by proceeding exactly as in the proofs of [Proposition 1](#) and [Theorem 1](#). \square

We stress here that [Proposition D.1](#) does not correspond to the Hedge instantiation [Theorem 1](#). Indeed, the second order term $\frac{1}{2} \sum_{t=1}^T \eta_t \mathbb{E}[G_t(\xi_t)^2 | \mathcal{F}_t]$ builds on results that are specific to Hedge, and is a priori considerably sharper than $\frac{\alpha^2}{2K} \sum_{t=1}^T \eta_t M_t^2$, the second order term of [Theorem 1](#).

D.3. Guarantees for Alg. 2. For clarity, we begin by reminding the specific assumptions relative to Alg. 2. In particular, we are still considering throughout a dual averaging policy (DA) with a negentropy regularizer. We additionally assume that at each round t , we receive a model u_t built according to the “smoothing” approach described in Section 5 where for all t :

$$u_t(x) = K_t(x_t, x) \cdot u_t(x_t) / p_t(x_t) \quad (\text{D.17})$$

where $K_t: \mathcal{K} \times \mathcal{K} \rightarrow \mathbb{R}$ is a (time-varying) *smoothing kernel*, i.e., $\int_{\mathcal{K}} K_t(x, x') dx' = 1$ for all $x \in \mathcal{K}$. For concreteness (and sampling efficiency), we will assume that payoffs now take values in $[0, 1]$, and we will focus on simple kernels that are supported on a neighborhood $\mathcal{U}_\delta(x) = \mathbb{B}_\delta(x) \cap \mathcal{K}$ of x in \mathcal{K} and are constant therein, i.e., $K^\delta(x, x') = [\lambda(\mathcal{U}_\delta(x))]^{-1} \mathbb{1}\{\|x' - x\| \leq \delta\}$. We incorporate in p_t an explicit exploration term of the form $\varepsilon_t / \lambda(\mathcal{K})$.

Under these assumptions, we may now bound both the bias and variance terms in (D.16).

Lemma D.6. *The following inequality holds, where L is a uniform Lipschitz coefficient for the reward functions u_t (as described in Assumption 1)*

$$B_t \leq L\delta_t. \quad (\text{D.18})$$

Moreover, there exists a constant $C_{\mathcal{K}}$ (depending only on the set \mathcal{K}) such that:

$$\sup_{\xi \in [0, 1]} \mathbb{E}[G_t(\xi)^2 | \mathcal{F}_t] \leq C_{\mathcal{K}} \delta_t^{-n} \epsilon_t^{-1}. \quad (\text{D.19})$$

Note that bounding the second order term of Theorem 1 under the same assumptions would have yielded a δ_t^{-2n} factor instead of δ_t^{-n} , which is a strictly weaker result!

Proof. We first prove (D.18). Using the fact that $u_t(x) = \int_{\mathcal{K}} u_t(x_t) K_t(x_t, x) dx_t$, we obtain:

$$\begin{aligned} |\mathbb{E}[u_t(x) - u_t(x_t) | \mathcal{F}_t]| &= \left| \int_{\mathcal{K}} K_t(x_t, x) \frac{u_t(x_t)}{p_t(x_t)} p_t(x_t) dx_t - \int_{\mathcal{K}} u_t(x) K_t(x_t, x) dx_t \right| \\ &= \left| \int_{\mathcal{K}} (u_t(x_t) - u_t(x)) K_t(x_t, x) dx_t \right| \\ &= [\lambda(\mathcal{U}_{\delta_t}(x))]^{-1} \left| \int_{\mathcal{K}} \mathbb{1}\{\|x' - x\| \leq \delta_t\} (u_t(x_t) - u_t(x)) dx_t \right| \\ &\leq L [\lambda(\mathcal{U}_{\delta_t}(x))]^{-1} \underbrace{\int_{\mathcal{K}} \mathbb{1}\{\|x' - x\| \leq \delta_t\} (u_t(x_t) - u_t(x)) dx_t}_{\leq \lambda(\mathcal{U}_{\delta_t}(x)) \delta_t} \leq L\delta_t. \end{aligned} \quad (\text{D.20})$$

This bound is uniform (does not depend on the point x), and thus implies the stated inequality for B_t .

We now turn to (D.19). To that end, let $\xi \in [0, 1]$. We will prove a uniform bound on $\mathbb{E}[G_t(\xi)^2 | \mathcal{F}_t]$. As a preliminary it is capital to note that, \mathcal{K} being convex compact, there exists constants $C_{\mathcal{K}}^M$ and $C_{\mathcal{K}}^m$ such that for all $x \in \mathcal{K}$,

$$C_{\mathcal{K}}^m \delta_t^n \leq [\lambda(\mathcal{U}_{\delta_t}(x))] \leq C_{\mathcal{K}}^M \delta_t^n.$$

Now, using $G_t(\xi)^2 = \langle \Lambda(\eta_t y_t + \xi \eta_t u_t), u_t^2 \rangle$ and $u_t(x) = K_t(x_t, x) \cdot u_t(x_t) / p_t(x_t)$, we may write:

$$\mathbb{E}[G_t(\xi)^2 | \mathcal{F}_t] = \mathbb{E} \left[\int_{\mathcal{K}} \Lambda(\eta_t y_t + \xi \eta_t u_t)(x) K_t(x_t, x)^2 \frac{u_t(x_t)^2}{p_t(x_t)^2} dx \middle| \mathcal{F}_t \right]$$

$$\begin{aligned}
&\leq \int_{\mathcal{K}} \overbrace{\frac{u_t(x')^2}{p_t(x')^2}}^{\leq 1} p_t(x') \left(\int_{\mathcal{K}} \Lambda(\eta_t y_t + \xi \eta_t u_t)(x) K_t(x', x)^2 dx \right) dx' \\
&\leq \int_{\mathcal{K}} \frac{1}{\underbrace{p_t(x')}_{\geq \varepsilon_t / \lambda(\mathcal{K})}} \underbrace{[\lambda(\mathcal{U}_{\delta_t}(x'))]^{-2}}_{\leq (C_{\mathcal{K}}^m)^2 \delta_t^{-2n}} \left(\int_{\mathcal{K}} \Lambda(\eta_t y_t + \xi \eta_t u_t)(x) \mathbb{1}\{\|x' - x\| \leq \delta_t\} dx \right) dx' \\
&\leq \frac{\lambda(\mathcal{K})}{(C_{\mathcal{K}}^m)^2 \delta_t^{2n} \varepsilon_t} \int_{\mathcal{K}} \left(\int_{\mathcal{K}} \Lambda(\eta_t y_t + \xi \eta_t u_t)(x) \mathbb{1}\{\|x' - x\| \leq \delta_t\} dx \right) dx' \\
&= \frac{\lambda(\mathcal{K})}{(C_{\mathcal{K}}^m)^2 \delta_t^{2n} \varepsilon_t} \int_{\mathcal{K}} \left(\int_{\mathcal{K}} \Lambda(\eta_t y_t + \xi \eta_t u_t)(x) \mathbb{1}\{\|x' - x\| \leq \delta_t\} dx' \right) dx \quad (\text{Fubini}) \\
&= \frac{\lambda(\mathcal{K})}{(C_{\mathcal{K}}^m)^2 \delta_t^{2n} \varepsilon_t} \int_{\mathcal{K}} \Lambda(\eta_t y_t + \xi \eta_t u_t)(x) \underbrace{\left(\int_{\mathcal{K}} \mathbb{1}\{\|x' - x\| \leq \delta_t\} dx' \right)}_{= \lambda(\mathcal{U}_{\delta_t}(x)) \leq C_{\mathcal{K}}^M \delta_t^n \lambda(\mathcal{K})} dx \\
&\leq \frac{\lambda(\mathcal{K})}{(C_{\mathcal{K}}^m)^2 \delta_t^{2n} \varepsilon_t} C_{\mathcal{K}}^M \delta_t^n \underbrace{\left(\int_{\mathcal{K}} \Lambda(\eta_t y_t + \xi \eta_t u_t)(x) dx \right)}_{=1} \\
&= \left(\frac{\lambda(\mathcal{K}) C_{\mathcal{K}}^M}{(C_{\mathcal{K}}^m)^2} \right) \delta_t^{-n} \varepsilon_t^{-1}. \tag{D.21}
\end{aligned}$$

This bound depends only on \mathcal{K} , and is notably independent on $\xi \in [0, 1]$. The result (D.19) follows directly. \square

We are now ready to prove [Proposition 2](#) and [Eq. \(23\)](#).

Proposition 2. *Suppose that the Hedge variant of [Alg. 2](#) is run with learning rate $\eta_t \propto 1/t^\rho$ and smoothing/exploration schedules $\delta_t \propto 1/t^\mu$, $\varepsilon_t \propto 1/t^\beta$ respectively. Then, the learner enjoys the bound*

$$\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^\rho + T^{1-\mu} + T^{1-\beta} + T^{1+n\mu+\beta-\rho}). \tag{22}$$

In particular, if the algorithm is run with $\rho = (n+2)/(n+3)$ and $\mu = \beta = 1/(n+3)$, we obtain the bound $\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^{\frac{n+2}{n+3}})$.

Proof. Let us consider a slight modification of [Alg. 2](#) in which

- The models (u_t) received by the learner are the same models than those generated by running [Alg. 2](#),
- At each round t , the action x_t is sampled according to $\tilde{p}_t = \Lambda(\eta_t y_t)$ (without taking into account the explicit exploration term).

The regret of this algorithm may be bounded using the Hedge template bound stated in [Proposition D.1](#), since we are indeed considering the regret induced by Hedge against the sequence of reward models u_t ⁴. Then, writing $\widetilde{\text{Reg}}(T)$ for the regret induced by the policy $(\tilde{p}_t)_t$, we get

$$\mathbb{E}[\widetilde{\text{Reg}}_x(T)] \leq \frac{\log(\lambda(\mathcal{K})/\lambda(\mathcal{C}))}{\eta_{T+1}} + L \text{diam}(\mathcal{C})T + 2 \sum_{t=1}^T B_t + \frac{1}{2} \sum_{t=1}^T \eta_t \mathbb{E}[G_t(\xi_t)^2 | \mathcal{F}_t]. \tag{D.22}$$

⁴Even though these models were generated by [Alg. 2](#), which does not exactly corresponds to Hedge

Using the bounds presented in [Lemma D.6](#) we then get:

$$\mathbb{E}[\widetilde{\text{Reg}}_x(T)] \leq \frac{\log(\lambda(\mathcal{K})/\lambda(\mathcal{C}))}{\eta_{T+1}} + L \text{diam}(\mathcal{C})T + 2L \sum_{t=1}^T \delta_t + \frac{1}{2}C_{\mathcal{K}} \sum_{t=1}^T \eta_t \delta_t^{-n} \epsilon_t^{-1}. \quad (\text{D.23})$$

We are however interested in guarantees for [Alg. 2](#), in which we play with the policy $(p_t)_t$, which slightly differs from the Hedge policy $(\tilde{p}_t)_t$. To that end, [Lemma D.3](#) enables us to bound the difference between the regrets $\text{Reg } T$ and $\widetilde{\text{Reg}}(T)$, induced by $(p_t)_t$ and $(\tilde{p}_t)_t$ respectively. Namely we can write:

$$\text{Reg } T \leq \widetilde{\text{Reg}}(T) + \sum_{t=1}^T \|\tilde{p}_t - p_t\|_{\infty}. \quad (\text{D.24})$$

For any $t \geq 1$, $x \in \mathcal{K}$, we have

$$|\tilde{p}_t(x) - p_t(x)| = \left| \tilde{p}_t(x) - (1 - \epsilon_t)\tilde{p}_t(x) - \frac{\epsilon_t}{\lambda(\mathcal{K})} \right| = \epsilon_t \cdot |\tilde{p}_t - 1/\lambda(\mathcal{K})| \leq \epsilon_t \left(1 + \frac{1}{\lambda(\mathcal{K})} \right). \quad (\text{D.25})$$

Injecting this in [\(D.24\)](#) we get

$$\text{Reg } T \leq \widetilde{\text{Reg}}(T) + \left(1 + \frac{1}{\lambda(\mathcal{K})} \right) \sum_{t=1}^T \epsilon_t. \quad (\text{D.26})$$

Finally, combining [\(D.26\)](#) with [\(D.23\)](#) yields:

$$\begin{aligned} \mathbb{E}[\widetilde{\text{Reg}}_x(T)] &\leq \frac{\log(\lambda(\mathcal{K})/\lambda(\mathcal{C}))}{\eta_{T+1}} + L \text{diam}(\mathcal{C})T \\ &\quad + 2L \sum_{t=1}^T \delta_t + \frac{C_{\mathcal{K}}}{2} \sum_{t=1}^T \eta_t \delta_t^{-n} \epsilon_t^{-1} + \left(1 + \frac{1}{\lambda(\mathcal{K})} \right) \sum_{t=1}^T \epsilon_t. \end{aligned} \quad (\text{D.27})$$

Now, using the same reasoning as in the proof of [Theorem 1](#) with regards to the set \mathcal{C} , and using $\eta_t \propto 1/t^\rho$, $\delta_t \propto 1/t^\mu$ and $\epsilon_t \propto 1/t^\beta$ straightforwardly gives:

$$\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^\rho + T^{1-\mu} + T^{1-\beta} + T^{1+n\mu+\beta-\rho}).$$

Finally, $\rho = (n+2)/(n+3)$ and $\mu = \beta = 1/(n+3)$ gives the optimal bound:

$$\mathbb{E}[\text{Reg}(T)] = \mathcal{O}(T^{\frac{n+2}{n+3}}).$$

□

Proposition 3. *Suppose that the Hedge variant of [Alg. 2](#) is run with parameters as in [Proposition 2](#) against a stream of loss functions with variation $V_T = \mathcal{O}(T^\nu)$. Then, the learner enjoys*

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1+n\mu+\beta-\rho} + T^{1-\beta} + T^{1-\mu} + T^{\nu+2\rho-n\mu-\beta}). \quad (23)$$

In particular, if the algorithm is run with $\rho = (1-\nu)(n+2)/(n+4)$ and $\mu = \beta = (1-\nu)/(n+4)$, we obtain the optimized bound $\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{\frac{n+3+\nu}{n+4}})$.

Proof. We use the same virtual segmentation as in the proof of [Theorem 2](#). As a reminder, this means that we partition the interval $\mathcal{T} = [1..T]$ into m contiguous segments \mathcal{T}_k , $k = 1, \dots, m$, each of length Δ (except possibly the m -th one, which might be smaller). More explicitly, take the window length to be of the form $\Delta = \lceil T^\gamma \rceil$ for some constant $\gamma \in [0, 1]$ to be determined later. In this way, the number of windows is $m = \lceil T/\Delta \rceil = \Theta(T^{1-\gamma})$ and the k -th window will be of the form $\mathcal{T}_k = [(k-1)\Delta + 1..k\Delta]$ for all $k = 1, \dots, m-1$ (the value $k = m$ is excluded as the m -th window might be smaller). For concision, we will denote the

learner’s static regret over the k -th window as $\text{Reg}(\mathcal{T}_k) = \max_{x \in \mathcal{K}} \sum_{t \in \mathcal{T}_k} \langle u_t, \delta_x - p_t \rangle$ (and likewise for its dynamic counterpart).

Following the proof of [Theorem 2](#) up to [\(C.30\)](#), we can still write in our bandit setting:

$$\text{DynReg}(T) \leq \sum_{k=1}^m \text{Reg}(\mathcal{T}_k) + 2\Delta V_T. \quad (\text{D.28})$$

Now [Proposition 2](#) applied to the Hedge variant of [Alg. 2](#) readily yields

$$\mathbb{E}[\text{Reg}(\mathcal{T}_k)] = \mathcal{O}\left((k\Delta)^\rho + \sum_{t \in \mathcal{T}_k} t^{-\beta} + \sum_{t \in \mathcal{T}_k} t^{-\mu} + \sum_{t \in \mathcal{T}_k} t^{\beta+n\mu-\rho}\right) \quad (\text{D.29})$$

so, after summing over all windows, we have

$$\begin{aligned} \sum_{k=1}^m \mathbb{E}[\text{Reg}(\mathcal{T}_k)] &= \mathcal{O}\left(\Delta^\rho \sum_{k=1}^m k^\rho + \sum_{t=1}^T t^{-\beta} + \sum_{t=1}^T t^{-\mu} + \sum_{t=1}^T t^{\beta+n\mu-\rho}\right) \\ &= \mathcal{O}(\Delta^\rho m^{1+\rho} + T^{1-\beta} + T^{1-\mu} + T^{1+\beta+n\mu-\rho}). \end{aligned} \quad (\text{D.30})$$

Since $\Delta = \mathcal{O}(T^\gamma)$ and $m = \mathcal{O}(T/\Delta) = \mathcal{O}(T^{1-\gamma})$, we get

$$\Delta^\rho m^{1+\rho} = \mathcal{O}((m\Delta)^\rho m) = \mathcal{O}(T^\gamma T^{(1-\gamma)(1+\rho)}) = \mathcal{O}(T^{1+\rho-\gamma}). \quad (\text{D.31})$$

Then, substituting in [\(D.30\)](#) and [\(D.28\)](#), we finally get the dynamic regret bound

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1+\rho-\gamma} + T^{1-\beta} + T^{1-\mu} + T^{1+\beta+n\mu-\rho} + T^\gamma V_T). \quad (\text{D.32})$$

To balance the above expression, we take $\gamma = 2\rho - \beta - n\mu$ for the window size exponent (which calibrates the first and fourth terms in the sum above). In this way, we finally obtain

$$\mathbb{E}[\text{DynReg}(T)] = \mathcal{O}(T^{1+n\mu+\beta-\rho} + T^{1-\beta} + T^{1-\mu} + T^{2\rho-n\mu-\beta} V_T) \quad (\text{D.33})$$

and our proof is complete. \square

APPENDIX E. NUMERICAL EXPERIMENTS

Our aim in this appendix is to provide some numerical illustrations of the theory presented in the rest of our paper. All numerical experiments were run on a machine with 48 CPUs (Intel(R) Xeon(R) Gold 6146 CPU @ 3.20GHz), with 2 Threads per core, and 500Go of RAM. For a simulation horizon of $T = 2 \times 10^5$, we choose a reward function $[0, 1]$ that is a linear combination of trigonometric terms with different frequencies and amplitudes, arbitrarily drawn. Because of this analytic expression, we are able to calculate the learner’s best action in hindsight (or instantaneously) and plot the relevant regret curves.

For illustration purposes, we compared 2 strategies, called “Grid” and “Kernel”. The “Kernel” method is as outlined in [Section 5](#) (cf. [Alg. 2](#)) with parameters described below. The “Grid” method involves partitioning the search space into a grid of a given mesh-size (a hyperparameter of the algorithm), and then treating the problem as a finite-armed bandit; in particular, the “Grid” strategy employs the EXP3 algorithm [\[7\]](#) with rewards sampled at the grid points.

In [Fig. 1](#), we plot the mean regret for both algorithms, with different hyperparameters, over T iterations. The confidence intervals are represented by the shaded areas, which corresponds to the mean value of the regret modulated by the standard deviation of our sample runs of each algorithm (computed on 92 initialization seeds for sampling, kept constant across different runs for control validation).

The dashed line represent in the figure corresponds to the theoretical regret bound of $T^{-\frac{1}{3}}$, which is the expected regret bound of the Kernel algorithm mean regret (without explicit exploration in our case). For performance evaluation purposes, we “slice” different snapshots




Figure 1: Expected average regret, averaged on 92 realizations for each algorithm (solid line). The variance is presented (shaded area) where we add and remove the standard deviation (computed on the 92 seeds) from the mean. Finally, the theoretical regret bound is displayed (dashed line).

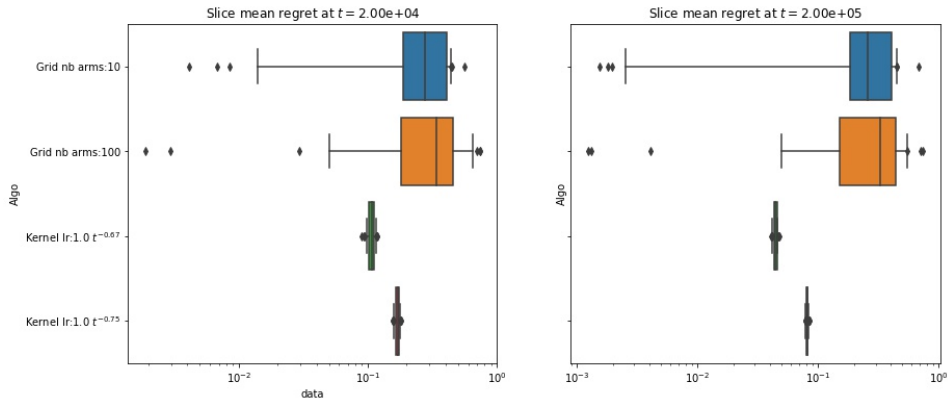


Figure 2: Two slices of the mean regret, averaged on 92 realizations for each algorithm (solid line). Whisker at 5-95% CI, boxes at 25-75% CI and median displayed with vertical bars.

of the regret in Fig. 2 at iteration counts 2×10^5 and 2×10^5 . In both cases, we observe a dramatic drop in variance for the Kernel algorithm relative to the Grid strategy, with a fixed number of arms uniformly cut beforehand; we also note that the performance of the Kernel method approaches the theoretical slope of $T^{-1/3}$ that characterizes the Kernel method.

By contrast, the mean regret for the Grid approach seems to converge to a finite value which indicates a much slower regret minimization rate; on the other hand, the mean regret of the Kernel method converges to 0 at the anticipated rate.

REFERENCES

[1] Jacob Abernethy, Peter L. Bartlett, Alexander Rakhlin, and Ambuj Tewari. Optimal strategies and minimax lower bounds for online convex games. In *COLT '08: Proceedings of the 21st Annual Conference*

- on *Learning Theory*, 2008.
- [2] Naman Agarwal, Alon Gonen, and Elad Hazan. Learning in non-convex games with an optimization oracle. In *COLT '19: Proceedings of the 32nd Annual Conference on Learning Theory*, 2019.
 - [3] Rajeev Agrawal. Sample mean based index policies with $o(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27(4):1054–1078, December 1995.
 - [4] Felipe Alvarez, Jérôme Bolte, and Olivier Brahic. Hessian Riemannian gradient flows in convex programming. *SIAM Journal on Control and Optimization*, 43(2):477–501, 2004.
 - [5] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
 - [6] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 1995.
 - [7] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
 - [8] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
 - [9] Balabhaskar Balasundaram, Sergiy Butenko, and Illya V. Hicks. Clique relaxations in social network analysis: The maximum k -plex problem. *Operations Research*, 59(1):133–142, January 2011.
 - [10] Claude Berge. *Topological Spaces*. Dover, New York, 1997.
 - [11] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244, October 2015.
 - [12] Avrim Blum and Adam Tauman Kalai. Universal portfolios with and without transaction costs. *Machine Learning*, 35(3):193–205, 1999.
 - [13] Immanuel M. Bomze, Marco Budinich, Panos M. Pardalos, and Marcello Pelillo. The maximum clique problem. In *Handbook of Combinatorial Optimization*. Springer, 1999.
 - [14] Mario Bravo and Panayotis Mertikopoulos. On the robustness of learning in games with stochastically perturbed payoff observations. *Games and Economic Behavior*, 103, John Nash Memorial issue:41–66, May 2017.
 - [15] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–358, 2015.
 - [16] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
 - [17] Sébastien Bubeck and Ronen Eldan. Multi-scale exploration of convex functions and bandit convex optimization. In *COLT '16: Proceedings of the 29th Annual Conference on Learning Theory*, 2016.
 - [18] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. \mathcal{X} -armed bandits. *Journal of Machine Learning Research*, 12:1655–1695, 2011.
 - [19] Sébastien Bubeck, Yin Tat Lee, and Ronen Eldan. Kernel-based methods for bandit convex optimization. In *STOC '17: Proceedings of the 49th annual ACM SIGACT symposium on the Theory of Computing*, 2017.
 - [20] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
 - [21] Nicolò Cesa-Bianchi, Pierre Gaillard, Gábor Lugosi, and Gilles Stoltz. Mirror descent meets fixed share (and feels no regret). In 989-997 (ed.), *Advances in Neural Information Processing Systems*, volume 25, 2012.
 - [22] Gong Chen and Marc Teboulle. Convergence analysis of a proximal-like minimization algorithm using Bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, August 1993.
 - [23] Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *COLT '12: Proceedings of the 25th Annual Conference on Learning Theory*, 2012.
 - [24] Benoît Duvocelle, Panayotis Mertikopoulos, Mathias Staudigl, and Dries Vermeulen. Learning in time-varying games. <https://arxiv.org/abs/1809.03066>, 2018.
 - [25] Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *SODA '05: Proceedings of the 16th annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 385–394, 2005.
 - [26] Gerald B. Folland. *Real Analysis*. Wiley-Interscience, 2 edition, 1999.

- [27] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.
- [28] Nadav Hallak, Panayotis Mertikopoulos, and Volkan Cevher. Regret minimization in stochastic non-convex learning via a proximal-gradient approach. <https://arxiv.org/abs/2010.06250>, 2020.
- [29] Elad Hazan. A survey: The convex optimization approach to regret minimization. In Suvrit Sra, Sebastian Nowozin, and Stephen J. Wright (eds.), *Optimization for Machine Learning*, pp. 287–304. MIT Press, 2012.
- [30] Elad Hazan and Comandur Seshadhri. Efficient learning algorithms for changing environments. In *ICML '09: Proceedings of the 26th International Conference on Machine Learning*, 2009.
- [31] Elad Hazan, Karan Singh, and Cyril Zhang. Efficient regret minimization in non-convex games. In *ICML '17: Proceedings of the 34th International Conference on Machine Learning*, 2017.
- [32] Mark Herbster and Manfred K. Warmuth. Tracking the best expert. *Machine Learning*, 32(2):151–178, 1998.
- [33] Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *AISTATS '15: Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015.
- [34] Adam Tauman Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, October 2005.
- [35] Narendra Karmarkar. Riemannian geometry underlying interior point methods for linear programming. In *Mathematical Developments Arising from Linear Programming*, number 114 in Contemporary Mathematics. American Mathematical Society, 1990.
- [36] Robert David Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *NIPS' 04: Proceedings of the 18th Annual Conference on Neural Information Processing Systems*, 2004.
- [37] Robert David Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-armed bandits in metric spaces. In *STOC '08: Proceedings of the 40th annual ACM symposium on the Theory of Computing*, 2008.
- [38] Walid Krichene, Maximilian Balandat, Claire Tomlin, and Alexandre Bayen. The Hedge algorithm on a continuum. In *ICML '15: Proceedings of the 32nd International Conference on Machine Learning*, 2015.
- [39] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- [40] Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2):465–507, January 2019.
- [41] Theodore S. Motzkin and Ernst G. Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canadian Journal of Mathematics*, 1965.
- [42] Yurii Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, 2009.
- [43] Steven Perkins, Panayotis Mertikopoulos, and David S. Leslie. Mixed-strategy learning with continuous action sets. *IEEE Trans. Autom. Control*, 62(1):379–384, January 2017.
- [44] Robert Ralph Phelps. *Convex Functions, Monotone Operators and Differentiability*. Lecture Notes in Mathematics. Springer-Verlag, 2 edition, 1993.
- [45] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [46] Shai Shalev-Shwartz and Yoram Singer. Convex repeated games and Fenchel duality. In *Advances in Neural Information Processing Systems 19*, pp. 1265–1272. MIT Press, 2007.
- [47] Aleksandrs Slivkins. Introduction to multi-armed bandits. *arXiv preprint arXiv:1904.07272*, 2019.
- [48] Aleksandrs Slivkins. Introduction to multi-armed bandits. *Foundations and Trends in Machine Learning*, 12(1-2):1–286, November 2019.
- [49] Victor Spirin and Leonid A. Mirny. Protein complexes and functional modules in molecular networks. *Proceedings of the National Academy of Sciences*, 2003.
- [50] Arun Sai Suggala and Praneeth Netrapalli. Online non-convex learning: Following the perturbed leader is optimal. In *ALT '20: Proceedings of the 31st International Conference on Algorithmic Learning Theory*, 2020.
- [51] Constantino Tsallis. Possible generalization of Boltzmann–Gibbs statistics. *Journal of Statistical Physics*, 52:479–487, 1988.
- [52] Paul Tseng. Convergence properties of Dikin’s affine scaling algorithm for nonconvex quadratic minimization. *Journal of Global Optimization*, 30(2):285–300, 2004.

- [53] Robert J. Vanderbei, Marc S. Mekeon, and Barry A. Freedman. A modification of Karmarkar's linear programming algorithm. *Algorithmica*, 1(1):395–407, November 1986.
- [54] Vladimir G. Vovk. Aggregating strategies. In *COLT '90: Proceedings of the 3rd Workshop on Computational Learning Theory*, pp. 371–383, 1990.
- [55] Lin Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11:2543–2596, October 2010.
- [56] Zhisheng Zhong, Tiancheng Shen, Yibo Yang, Chao Zhang, and Zhouchen Lin. Joint sub-bands learning with clique structures for wavelet domain super-resolution. In *NeurIPS '18: Proceedings of the 32nd International Conference of Neural Information Processing Systems*, 2018.
- [57] Zhengyuan Zhou, Panayotis Mertikopoulos, Nicholas Bambos, Stephen P. Boyd, and Peter W. Glynn. On the convergence of mirror descent beyond stochastic convex programming. *SIAM Journal on Optimization*, 30(1):687–716, 2020.
- [58] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML '03: Proceedings of the 20th International Conference on Machine Learning*, pp. 928–936, 2003.