

Refined approachability algorithms and application to regret minimization with global costs

Joon Kwon

*MIA Paris, INRAE & AgroParisTech
16 rue Claude Bernard, 75231 Paris, France*

JOON.KWON@INRAE.FR

Editor: Gábor Lugosi

Abstract

Blackwell’s approachability is a framework where two players, the Decision Maker and the Environment, play a repeated game with vector-valued payoffs. The goal of the Decision Maker is to make the average payoff converge to a given set called the target. When this is indeed possible, simple algorithms which guarantee the convergence are known. This abstract tool was successfully used for the construction of optimal strategies in various repeated games, but also found several applications in online learning. By extending an approach proposed by Abernethy et al. (2011), we construct and analyze a class of Follow the Regularized Leader algorithms (FTRL) for Blackwell’s approachability which are able to minimize not only the Euclidean distance to the target set (as it is often the case in the context of Blackwell’s approachability) but a wide range of distance-like quantities. This flexibility enables us to apply these algorithms to closely minimize the quantity of interest in various online learning problems. In particular, for regret minimization with ℓ_p global costs, we obtain the first bounds with explicit dependence in p and the dimension d .

Keywords: Blackwell’s Approachability, Follow the Regularized Leader, Online Learning, Regret Minimization, Global Costs

1. Introduction

One of the foundational results of game theory is von Neumann’s minimax theorem which characterizes the highest payoff that each player of a finite zero-sum game can guarantee regardless of the opponent’s strategy. In the seminal works of Blackwell (1956, 1954), a surprising extension of this result was proposed in the context of repeated games with vector-valued payoffs. The so-called Blackwell’s condition characterizes the convex sets that the player can guarantee to asymptotically reach, regardless of the opponent’s actions. In the case of non-convex sets, this condition remains sufficient. When the above condition is satisfied for a given set called the *target*, the original algorithm proposed by Blackwell guarantees that the average vector-valued payoff converges to (*approaches*) the target set at rate $O(1/\sqrt{T})$, where T is the number of rounds of the repeated play. This topic is now called Blackwell’s approachability.

This framework was used for the construction of optimal strategies in repeated games as in Kohlberg (1975), see also the survey work by Perchet (2014) and references therein. Beyond the field of game theory, this tool has been noticed by the machine learning community and used for constructing and analyzing algorithms for various online decision problems such as regret minimization (Cesa-Bianchi and Lugosi, 2006), asymptotic calibration

(Dawid, 1982; Foster and Vohra, 1998), regret minimization with variable stage duration (Mannor and Shimkin, 2008) or with global cost functions (Even-Dar et al., 2009). However, one drawback of using Blackwell’s approachability is that algorithms then usually minimize the *Euclidean distance* of the average payoffs to the target set, which is seldom the exact quantity of interest in online learning applications. One of the main objectives of the present work is to provide a flexible class of algorithms which are able to minimize various distance-like quantities, and not only the Euclidean distance.

Several alternative approachability algorithms were also proposed, including potential-based algorithm (Hart and Mas-Colell, 2001) which generalize the Euclidean projection involved in Blackwell’s algorithm, and response-based algorithms (Bernstein and Shimkin, 2015) which avoid the projection altogether. Besides, an important scheme used in several works is the conversion of regret minimization algorithms into approachability algorithms (Abernethy et al., 2011; Shimkin, 2016; Mannor et al., 2014).

Regret minimization was introduced by Hannan (1957) and is a sequential decision problem where the Decision Maker aims at minimizing the difference between its payoff and the highest payoff in hindsight given by a constant strategy. The link between approachability and regret minimization was already noticed by Blackwell (1954) who reduced regret minimization to an approachability problem. Hart and Mas-Colell (2001) proposed an alternative reduction and constructed a whole family of regret minimization algorithm using potential-based approachability algorithms. Gordon (2007) extended the potential-based approach to a wider range of regret minimization problems, seen as approachability problems. Conversely, regret minimizing algorithms have been converted into approachability algorithms (Gordon, 2007; Abernethy et al., 2011; Perchet, 2015; Shimkin, 2016).

It is worth noting that modern variants of the Regret Matching algorithm, which is a special case of potential-based approachability algorithms (Hart and Mas-Colell, 2000, 2001), are today the state-of-the-art online learning algorithms for Nash equilibrium computation in large zero-sum games (Zinkevich et al., 2007; Tammelin et al., 2015).

1.1 Related work

In Perchet (2015), the Exponential Weights Algorithm, which is a central regret minimization algorithm, is adapted to approachability, and the resulting algorithm minimizes the ℓ_∞ distance to the target set.

The conversion scheme presented in Abernethy et al. (2011) deals with online linear optimization algorithms which are transposed into the approachability of convex cone target sets, and the associated guarantee is an upper bound on the Euclidean distance to the target set. An extension to all convex target sets is also given, which involves the adding of a dimension.

A closely related work is Shimkin (2016) where a conversion from online *convex* optimization algorithm to approachability of *bounded* convex sets is presented, which guarantees an upper bound on the distance to the target set measured with the Euclidean norm or possibly any other norm.

One of the applications of approachability is the problem of regret minimization with global costs, introduced in Even-Dar et al. (2009) and already analyzed as an approachability problem. This problem was further studied in Rakhlin et al. (2011); Bernstein and Shimkin

(2015), and in a recent paper (Liu et al., 2021), the authors used the conversion scheme from Shimkin (2016) to construct and analyze algorithms for this problem.

A recent paper (Farina et al., 2021) proposes an extension of Abernethy et al. (2011) by introducing predictive approachability algorithms. The proposed construction shares similarities with the present work but focuses on variants of Regret Matching, and only derives upper bounds on Euclidean distances to target sets—see (Farina et al., 2021, Proposition 2), whereas we consider a more general class of quantities in Definition 3 and Theorem 10 below.

1.2 Contributions

- We consider a class of Follow the Regularized Leader algorithms (FTRL) which we convert from regret minimization to approachability. The conversion scheme we use is a refinement of Abernethy et al. (2011), which itself is an extension of Gordon (2007), and the algorithms that we obtain are capable of minimizing not only the Euclidean distance to the target set as in Abernethy et al. (2011), but the distance measured by an arbitrary norm, or even more general distance-like quantities. This flexibility will prove itself useful in the construction of tailored algorithms with tight bounds for various problems.
- For the problem of regret minimization with global cost, we construct algorithms for arbitrary norm cost functions and obtain novel guarantees. In particular, for ℓ_p norm cost functions ($p > 1$), we obtain the first explicit regret bounds that depend on p and the dimension d , and which recovers, in the special case $p = \infty$ the best known $O(\sqrt{(\log d)/T})$ bound.

1.3 Summary

In Section 2, we present a model of approachability with target sets which are closed convex cones. In Section 3, we define a class of FTRL algorithms and derive a general guarantee. In Section 4, we recall the problem of regret minimization with global cost functions and relate it to our approachability framework and FTRL algorithms. In the special case of ℓ_p norm cost functions, we derive regret bounds with explicit dependence in d and p . In Appendix D, we recall Blackwell’s algorithm and prove that it belongs to the class of algorithms defined in Section 3. In Appendix E, we present a variant of the model from Section 2, where the Decision Maker may choose its actions at random from a finite set. We then define corresponding FTRL algorithms and provide guarantees in expectation, with high probability and almost-surely. In Appendices G and H, we recall the problems of online combinatorial optimization and internal/swap regret respectively, their reductions to approachability problems, and demonstrate that a carefully chosen FTRL algorithm recover the known optimal bounds.

1.4 Notation

\mathbb{R}_+^* denotes the set of positive real numbers. $d \geq 2$ will always denote an integer. All vector spaces will be of finite dimension. For $p \in [1, +\infty]$, we denote $\|\cdot\|_p$ the ℓ_p norm, meaning for

$x \in \mathbb{R}^d$, $\|x\|_p = \left(\sum_{i=1}^d |x_i|^p\right)^{1/p}$ for $p < +\infty$ and $\|x\|_\infty = \max_{1 \leq i \leq d} |x_i|$. For a given norm $\|\cdot\|$ in a vector space, the dual norm $\|\cdot\|_*$ is defined by $\|y\|_* = \sup_{\|x\| \leq 1} |\langle y, x \rangle|$. Denote Δ_d the unit simplex of \mathbb{R}^d : $\Delta_d = \left\{x \in \mathbb{R}_+^d, \sum_{i=1}^d x_i = 1\right\}$. For a sequence $(r_t)_{t \geq 1}$ of vectors, we denote $\bar{r}_T = \frac{1}{T} \sum_{t=1}^T r_t$ the average of the T first terms ($T \geq 1$). If \mathcal{X} a subset of a vector space, $I_{\mathcal{X}}$ denotes the convex indicator of \mathcal{X} , in other words: $I_{\mathcal{X}}(x) = 0$ if $x \in \mathcal{X}$ and $I_{\mathcal{X}}(x) = +\infty$ otherwise. If a vector $x_t \in \mathbb{R}^d$ is denoted with an index (t in this example), its components are denoted with an additional index as follows: $x_t = (x_{ti})_{1 \leq i \leq d}$.

2. Approachability of convex cones

We introduce a simple repeated game with vector-valued payoffs between two players (the Decision Maker and the Environment) with a closed convex cone target set for the Decision Maker. We then state a few properties about closed convex cones and support functions.

2.1 Model

Let \mathcal{V} be a finite-dimensional vector space and denote \mathcal{V}^* its dual. The latter will be the *payoff space*. Let \mathcal{A} and \mathcal{B} be the *action sets* for the Decision Maker and the Environment respectively, about which we assume no special structure. Let $r: \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{V}^*$ be a vector-valued *payoff function*. The game is played as follows. At time $t \geq 1$,

- the Decision Maker chooses action $a_t \in \mathcal{A}$;
- the Environment chooses action $b_t \in \mathcal{B}$;
- the Decision Maker observes *vector payoff* $r_t := r(a_t, b_t) \in \mathcal{V}^*$.

We allow the Environment to be *adversarial*¹.

The problem involves a *target set* $\mathcal{C} \subset \mathcal{V}^*$ which we assume to be a closed convex cone². The goal is to construct algorithms which guarantee that the average payoff $\bar{r}_T := \frac{1}{T} \sum_{t=1}^T r_t$ is *close* to the target \mathcal{C} in a sense that will be made precise.

The above model does not allow the Decision Maker to choose actions at random. Such a model is presented in Appendix E.

2.2 Generator of a closed convex cone

We now introduce a key notion of this work which will be used in Section 2.3 to define the class of quantities that will be minimized by the algorithms defined in Section 3.2. Definitions and properties about closed convex cones are gathered in Appendix A.

Definition 1 *Let \mathcal{C} be a closed convex cone. A set \mathcal{X} is a generator of \mathcal{C} if it is convex, compact and if $\mathbb{R}_+ \mathcal{X} = \mathcal{C}$.*

1. In other words, action b_t chosen by the Environment may depend on anything that has happened before it is chosen, including a_t .

2. For the case where target set is a closed convex set but not a cone, we refer to (Abernethy et al., 2011, Section 4 & Lemma 14) where a conversion scheme into an auxiliary problem where the target is a cone is presented.

The following proposition gives three examples of generators. The second example demonstrates that a generator always exists. The proof is given in Appendix A.1.

Proposition 2 *Let \mathcal{W} be the ambient finite-dimensional vector space.*

- (i) *If $\mathcal{W} = \mathcal{W}^* = \mathbb{R}^d$, the negative orthant \mathbb{R}_-^d is a closed convex cone and $(\mathbb{R}_-^d)^\circ = \mathbb{R}_+^d$. Moreover, Δ_d is a generator of \mathbb{R}_+^d .*
- (ii) *Let $\mathcal{C} \subset \mathcal{W}$ be a closed convex cone, $\|\cdot\|$ a norm on \mathcal{W} , and \mathcal{B} the closed unit ball with respect to $\|\cdot\|$. Then, $\mathcal{B} \cap \mathcal{C}$ is a generator of \mathcal{C} .*
- (iii) *If \mathcal{X} is a nonempty convex compact subset of \mathcal{W} , then \mathcal{X} is a generator of $\mathcal{X}^{\circ\circ} = \mathbb{R}_+ \mathcal{X}$.*

2.3 Support functions

We now present support functions which will be used in Section 3.2 to express the quantities that will be minimized by our algorithms.

Definition 3 *For a nonempty subset $\mathcal{X} \subset \mathcal{V}$, the application $I_{\mathcal{X}}^* : \mathcal{V}^* \rightarrow \mathbb{R} \cup \{+\infty\}$ defined by*

$$I_{\mathcal{X}}^*(y) = \sup_{x \in \mathcal{X}} \langle y, x \rangle, \quad y \in \mathcal{V}^*,$$

is called the support function of \mathcal{X} .

The support function can be written as the Legendre–Fenchel transform of the indicator function of the set \mathcal{X} . It is therefore convex. Moreover, in the case where \mathcal{X} is a generator of the polar cone \mathcal{C}° of some closed convex cone $\mathcal{C} \subset \mathcal{V}^*$, the properties of $I_{\mathcal{X}}^*$ make it suitable for measuring how far a point of \mathcal{V}^* is from \mathcal{C} . Indeed, it is easy to check that $I_{\mathcal{X}}^*$ is then real-valued, continuous, and that for all points $y \in \mathcal{V}^*$,

$$I_{\mathcal{X}}^*(y) \leq 0 \quad \iff \quad y \in \mathcal{C}.$$

The following proposition demonstrates that the distance to a closed convex cone \mathcal{C} with respect to an arbitrary norm can be written as a support function. It is an extension of Lemma 13 in Abernethy et al. (2011) to an arbitrary norm. The proof is given in Appendix C.1.

Proposition 4 *Let \mathcal{C} be a closed convex cone in \mathcal{V}^* , $\|\cdot\|$ a norm on \mathcal{V} and $\|\cdot\|_*$ its dual norm on \mathcal{V}^* . Then,*

$$\inf_{y' \in \mathcal{C}} \|y' - y\|_* = I_{\mathcal{B} \cap \mathcal{C}^\circ}^*(y), \quad y \in \mathcal{V}^*,$$

where \mathcal{B} is the closed unit ball for $\|\cdot\|$.

2.4 Blackwell’s condition

In the case of convex sets, Blackwell’s condition (Blackwell, 1956) is a characterization of the target sets to which the Decision Maker can guarantee a convergence. We here present the special case of convex cones, which will be used in the construction and the analysis of the algorithms in Section 3.2.

Definition 5 (Blackwell’s condition for convex cones) *A closed convex cone \mathcal{C} of the payoff space \mathcal{V}^* is a B-set for the game $(\mathcal{A}, \mathcal{B}, r)$ if*

$$\forall x \in \mathcal{C}^\circ, \exists a(x) \in \mathcal{A}, \forall b \in \mathcal{B}, \quad \langle r(a(x), b), x \rangle \leq 0.$$

Such an application $a : \mathcal{C}^\circ \rightarrow \mathcal{A}$ is called a $(\mathcal{A}, \mathcal{B}, r, \mathcal{C})$ -oracle.

The geometric interpretation of this condition is that for any given hyperplane containing the target, the Decision Maker has an action which forces the payoff vector to belong the same side of the hyperplane as the target, regardless of the Environment’s action.

In some situations, it is easier to establish the following equivalent dual condition. The proof is given in Appendix C.2 for completeness.

Proposition 6 (Blackwell’s dual condition) *We assume that \mathcal{A}, \mathcal{B} are convex sets of finite dimensional vectors spaces, such that \mathcal{A} is compact, and that the payoff function $r : \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{V}^*$ is bi-affine. Then, a closed convex cone \mathcal{C} of the payoff space \mathcal{V}^* is a B-set for the game $(\mathcal{A}, \mathcal{B}, r)$ if, and only if*

$$\forall b \in \mathcal{B}, \exists a \in \mathcal{A}, \quad r(a, b) \in \mathcal{C}.$$

3. A class of FTRL algorithms

We define a class of Follow the Regularized Leader algorithms (FTRL) which are transposed from regret minimization, and which guarantee, when the target is a B-set, that the average payoff converges to the target set, the convergence being measured in a sense that will be made precise.

3.1 Regularizers

We first introduce regularizers functions and the notion of strong convexity needed for the definition and the analysis of FTRL algorithms (Shalev-Shwartz, 2007, 2011; Bubeck, 2011), which are also known as *dual averaging* (Nesterov, 2009) in the context of optimization. These are classic: basic properties, proofs and important examples are recalled in Appendix B. Again, \mathcal{V} and \mathcal{V}^* are finite-dimensional vectors spaces and \mathcal{X} is a nonempty convex compact subset of \mathcal{V} . We recall that the *domain* $\text{dom } h$ of a function $h : \mathcal{V} \rightarrow \mathbb{R} \cup \{+\infty\}$ is the set of points where it has finite values.

Definition 7 *A convex function $h : \mathcal{V} \rightarrow \mathbb{R} \cup \{+\infty\}$ is a regularizer on \mathcal{X} if it is strictly convex, lower semicontinuous, and has \mathcal{X} as domain.*

Definition 8 *Let $h : \mathcal{V} \rightarrow \mathbb{R} \cup \{+\infty\}$ be a function, $\|\cdot\|$ a norm on \mathcal{V} , and $K > 0$. h is K -strongly convex with respect to $\|\cdot\|$ if for all $x, x' \in \mathcal{V}$ and $\lambda \in [0, 1]$,*

$$h(\lambda x + (1 - \lambda)x') \leq \lambda h(x) + (1 - \lambda)h(x') - \frac{K\lambda(1 - \lambda)}{2} \|x' - x\|^2. \quad (1)$$

3.2 Definition and analysis of the algorithm

We now construct the FTRL algorithms for the model introduced in Section 2.1 and establish guarantees.

Let \mathcal{C} be a B-set for the game $(\mathcal{A}, \mathcal{B}, r)$ and $a : \mathcal{C}^\circ \rightarrow \mathcal{A}$ a $(\mathcal{A}, \mathcal{B}, r, \mathcal{C})$ -oracle. Let $\mathcal{X} \subset \mathcal{V}$ be a generator of \mathcal{C}° , h a regularizer on \mathcal{X} , and $(\eta_t)_{t \geq 1}$ a positive sequence of parameters. The associated algorithm is then defined for $t \geq 1$ as:

$$\begin{aligned} \text{compute } & x_t = \arg \max_{x \in \mathcal{X}} \left\{ \left\langle \eta_{t-1} \sum_{s=1}^{t-1} r_s, x \right\rangle - h(x) \right\} \\ \text{compute } & a_t = a(x_t) \\ \text{observe } & r_t = r(a_t, b_t), \end{aligned}$$

where the first line is well-defined thanks to the basic properties of regularizers gathered in Proposition 20. We prove in Appendix D that Blackwell's original algorithm belongs to this class.

The above definition of x_t can be interpreted as the action played by a FTRL algorithm in an online linear optimization problem with action set \mathcal{X} and payoff vectors $(r_t)_{t \geq 1}$. We state in the following lemma the classical *regret bound* guaranteed by such an algorithm (Shalev-Shwartz, 2007, 2011; Bubeck, 2011). The proof is given in Appendix C.3 for completeness. This regret bound will then be *converted* in Theorem 10 into an upper bound on $I_{\mathcal{X}}^*(\bar{r}_T)$, thus providing a guarantee for the approachability game. This conversion is an extension of the scheme introduced in Abernethy et al. (2011), which gives approachability algorithms which minimize the Euclidean distance of the average payoff to the target set. Our approach is more general as it allows, by the choice of the generator \mathcal{X} , to minimize a whole class of distance-like quantities.

The conversion is here applied to FTRL algorithms, but could have been applied to any online linear optimization algorithm.

In a recent paper (Farina et al., 2021), the authors also propose a similar extension of Abernethy et al. (2011) which is however less general, as they only consider generators which contain $\mathcal{C}^\circ \cap \mathcal{B}_2$, where \mathcal{B}_2 is the Euclidean ball.

Lemma 9 (Regret bound) *Let $\Delta, K, M > 0$, $\|\cdot\|$ a norm on \mathcal{V} , and $\|\cdot\|_*$ its dual norm on \mathcal{V}^* . We assume:*

- (i) $\max_{x \in \mathcal{X}} h(x) - \min_{x \in \mathcal{X}} h(x) \leq \Delta$,
- (ii) h is K -strongly convex with respect to $\|\cdot\|$,
- (iii) $\|r_t\|_* \leq M$ for all $t \geq 1$.

Then, the choice $\eta_t = \sqrt{\Delta K / M^2 t}$ (for $t \geq 1$) guarantees

$$\forall T \geq 1, \quad \max_{x \in \mathcal{X}} \sum_{t=1}^T \langle r_t, x \rangle - \sum_{t=1}^T \langle r_t, x_t \rangle \leq 2M \sqrt{\frac{\Delta T}{K}}.$$

The following theorem provides upper bounds on $I_{\mathcal{X}}^*(\bar{r}_T)$ (where $\bar{r}_T = \frac{1}{T} \sum_{t=1}^T r_t$ is the average payoff) and not only the Euclidean distance from \bar{r}_T to \mathcal{C} , which is a special case—see Proposition 4. Therefore, the choice of the generator \mathcal{X} determines the quantity that is minimized by the algorithm. We present in Sections 4 and Appendices G and H examples of problems where a judicious choice of generator \mathcal{X} allows $I_{\mathcal{X}}^*(\bar{r}_T)$ to be equal (or close) to the quantity the Decision Maker actually aims at minimizing and therefore yields tailored algorithms.

Theorem 10 *Let $\Delta, K, M > 0$, $\|\cdot\|$ a norm on \mathcal{V} , and $\|\cdot\|_*$ its dual norm on \mathcal{V}^* . We assume:*

- (i) $\max_{x \in \mathcal{X}} h(x) - \min_{x \in \mathcal{X}} h(x) \leq \Delta$,
- (ii) h is K -strongly convex with respect to $\|\cdot\|$,
- (iii) $\|r(a, b)\|_* \leq M$ for all $a \in \mathcal{A}$ and $b \in \mathcal{B}$.

Then the above algorithm guarantees, with the choice $\eta_t = \sqrt{\Delta K / M^2 t}$ (for $t \geq 1$), against any sequence of actions $(b_t)_{t \geq 1}$ chosen by the Environment,

$$\forall T \geq 1, \quad I_{\mathcal{X}}^*(\bar{r}_T) \leq 2M \sqrt{\frac{\Delta}{KT}}.$$

Proof The regret from Lemma 9 is the following quantity:

$$\text{Reg}_T = \max_{x \in \mathcal{X}} \sum_{t=1}^T \langle r_t, x \rangle - \sum_{t=1}^T \langle r_t, x_t \rangle.$$

The first term above can be written

$$\max_{x \in \mathcal{X}} \sum_{t=1}^T \langle r_t, x \rangle = T \cdot \max_{x \in \mathcal{X}} \left\langle \frac{1}{T} \sum_{t=1}^T r_t, x \right\rangle = T \cdot I_{\mathcal{X}}^*(\bar{r}_T),$$

whereas the second sum is nonpositive because each term is. Indeed, by definition of the algorithm, and because a is a $(\mathcal{A}, \mathcal{B}, r, \mathcal{C})$ -oracle,

$$\langle r_t, x_t \rangle = \langle r(a_t, b_t), x_t \rangle = \langle r(a(x_t), b_t), x_t \rangle \leq 0.$$

Therefore $I_{\mathcal{X}}^*(\bar{r}_T) \leq \frac{1}{T} \text{Reg}_T$ and the regret bound from Lemma 9 gives the result. ■

In Appendix E, we present a variant of the present model where the Decision Maker can choose its actions at random. The above guarantee is transposed into guarantees in expectation, in high-probability (using the Azuma–Hoeffding inequality), and into almost-sure convergence (using a Borel–Cantelli argument).

4. Regret minimization with global costs

The problem of regret minimization with global costs was introduced in [Even-Dar et al. \(2009\)](#). It is an adversarial online learning problem motivated by load balancing and job scheduling, where at each step, the Decision Maker first chooses a distribution (task allocation) over d machines, and then observes the cost of using each machine, which may be different for each machine and each step. The goal of the Decision Maker is to minimize, not the sum of the cumulative costs of using each machine, but a given function of the vector of cumulative costs. A typical example of such *global cost* function is the ℓ_p norm, which includes as special cases the sum of the costs (for $p = 1$), as well as the makespan i.e. the highest cumulative cost (for $p = \infty$). A very common approach for this type of problem is to focus on competitive ratio ([Borodin and El-Yaniv, 1998](#); [Azar et al., 1993](#); [Molinaro, 2017](#)). We instead follow [Even-Dar et al. \(2009\)](#) and aim at minimizing the regret.

In the seminal paper by [Even-Dar et al. \(2009\)](#), the authors introduce a reduction of the problem to an approachability game and obtain a regret bound of order $O((\log d)/\sqrt{T})$ for the ℓ_∞ cost function. For general convex cost functions, the authors present a regret bound that reads $\sqrt{d/T}$; however, this expression does not reflect the true dependency of the bound in the number d of machines, as this bound also involves several Lipschitz constants that depend on the cost function, and which may also depend on d , as it is the case for ℓ_p cost functions. In a theoretical work, [Rakhlin et al. \(2011\)](#) proved that the regret bound can be improved to $O(\sqrt{(\log d)/T})$ in the ℓ_∞ case, but no algorithm achieving this bound was provided. [Bernstein and Shimkin \(2015\)](#) also studied alternative algorithms for minimizing regret with global cost but no explicit bound was given. In a recent paper by [Liu et al. \(2021\)](#), new algorithms are proposed, based on a technique for adapting online convex optimization algorithms to approachability games ([Shimkin, 2016](#)), and regret bounds for monotone norms cost functions (which include ℓ_p norms) are derived. The bounds are abstract, except for the ℓ_∞ case where the algorithm achieve the best known $O(\sqrt{(\log d)/T})$ bound in addition of being the first such algorithm to run in polynomial time. Besides, more general problems than the one we consider below are studied in [Azar et al. \(2014\)](#); [Mannor et al. \(2014\)](#) and both provide algorithms with convergence rate $T^{-1/4}$.

In this section, we apply the tools introduced in Sections 2 and 3 to construct and analyze new algorithms for this problem. Although our approach applies to general norm cost functions (unlike [Liu et al. \(2021\)](#) which assumes the norm to be monotone), we focus in Section 4.4 on ℓ_p norms ($p > 1$) to obtain explicit regret bounds in Theorem 14, which, in the special case $p = \infty$, recovers the best known $O(\sqrt{(\log d)/T})$ bound. To the best of our knowledge, these are the first regret bounds for ℓ_p norm cost functions with explicit dependence in d and p .

We use the reduction of the problem to an approachability game from [Even-Dar et al. \(2009\)](#). We then choose a generator of the polar of the target set based on a specially crafted norm on the payoff space, which then enables us to bound the regret with cost functions by a support function. Then, in the case of ℓ_p cost functions, the explicit regret bounds are derived with the help of a carefully chosen regularizer.

4.1 Problem statement

Let $d \geq 2$ be an integer and $\|\cdot\|$ a norm on \mathbb{R}^d . Recall that Δ_d denotes the unit simplex of \mathbb{R}^d and is identified with the set of probability distributions over \mathcal{I} . For $t \geq 1$,

- the Decision Maker chooses distribution $a_t \in \Delta_d$;
- the Environment chooses loss vector $\ell_t \in [0, 1]^d$.

The Decision Maker aims at minimizing the following average regret:

$$\overline{\text{Reg}}_T = \left\| \frac{1}{T} \sum_{t=1}^T a_t \odot \ell_t \right\| - \min_{a \in \Delta_d} \left\| \frac{1}{T} \sum_{t=1}^T a \odot \ell_t \right\|,$$

where \odot denotes the component-wise multiplication. At each stage $t \geq 1$, the i -th component of vector $a_t \odot \ell$ is equal to $a_{ti}\ell_i$ and corresponds to the cost of using machine i for a fraction a_{ti} of the job. The regret is the difference between the actual global cost incurred by the Decision Maker and the best possible global cost in hindsight for a static distribution $a \in \Delta_d$. Important special cases include the makespan which corresponds to $\|\cdot\| = \|\cdot\|_\infty$: the global cost is then the highest average cost over the machines; and for $\|\cdot\| = \|\cdot\|_1$ the global cost simply corresponds to the sum of the costs of all the machines, and the problem then reduces to basic regret minimization.

4.2 Reduction to an approachability game

We recall the reduction given in (Even-Dar et al., 2009, Section 4) of the above problem to an approachability game which fits the model from Section 2.

Consider the following action sets for the Decision Maker and the Environment respectively: $\mathcal{A} = \Delta_d$ and $\mathcal{B} = [0, 1]^d$. Define the payoff function $r : \Delta_d \times [0, 1]^d \rightarrow (\mathbb{R}^d)^2$ as

$$r(a, \ell) = (a \odot \ell, \ell), \quad a \in \Delta_d, \ell \in [0, 1]^d,$$

and consider the following target set:

$$\mathcal{C} = \left\{ (y, y') \in (\mathbb{R}_+^d)^2, \|y\| \leq \min_{a \in \Delta_d} \|a \odot y'\| \right\}.$$

The payoff space is therefore $\mathcal{V}^* = (\mathbb{R}^d)^2$.

Proposition 11 (Even-Dar et al. (2009)) \mathcal{C} is a closed convex cone. Moreover, it is a B -set for the game $(\Delta_d, [0, 1]^d, r)$.

Proof We give the proof for the sake of completeness and essentially follow (Even-Dar et al., 2009, Lemma 5 & Theorem 6). \mathcal{C} can be written as

$$\mathcal{C} = \left\{ (y, y') \in (\mathbb{R}_+^d)^2 \mid \|y\| - \min_{a \in \Delta_d} \|a \odot y'\| \leq 0 \right\},$$

which then appears as a closed level set of a convex function because $y \mapsto \|y\|$ is continuous and convex for all norms, and because $y' \mapsto \min_{a \in \Delta_d} \|a \odot y'\|$ is concave on \mathbb{R}_+^d according

to (Rakhlin et al., 2011, Lemma 22) and continuous as the minimum of a family of continuous functions. \mathcal{C} is thus closed and convex, and because it is clearly closed by multiplication by a nonnegative scalar, it is a closed convex cone.

We can now establish that \mathcal{C} is a B-set for the game $(\Delta_d, [0, 1]^d, r)$ using Blackwell's dual condition from Proposition 6, because the payoff function r is indeed bi-affine. Let $\ell \in [0, 1]^d$ and consider $a_0 = \arg \min_{a \in \Delta_d} \|a \odot \ell\|$. Then, we clearly have $r(a_0, \ell) \in \mathcal{C}$, which concludes the proof. \blacksquare

Remark 12 (Computation of the oracle) *As noted in (Even-Dar et al., 2009, Section 4) and (Liu et al., 2021, Section 4.1), a $(\Delta_d, [0, 1]^d, r, \mathcal{C})$ -oracle is given by*

$$a(z, z') = \arg \min_{a \in \Delta_d} \sum_{i=1}^d \max(0, z_i a_i + z'_i), \quad (z, z') \in \mathcal{C}^\circ,$$

which is a linear program with $O(d)$ variables and $O(d)$ constraints, which can thus be computed in polynomial time.

4.3 A special norm on the payoff space

We now define a special norm on the payoff space \mathcal{V}^* which will allow us to bound the regret from above with the help of a support function, and will therefore provide the generator of \mathcal{C}° for defining the regularizer and constructing our algorithm.

We introduce the following norm $\|\cdot\|_{\mathcal{V}^*}$ whose definition is based on the norm $\|\cdot\|$ given in Section 4.1:

$$\|(y, y')\|_{\mathcal{V}^*} = \|y\| + \max_{a \in \Delta_d} \|a \odot y'\|, \quad (y, y') \in \mathcal{V}^* = (\mathbb{R}^d)^2.$$

It is easy to check that $\|\cdot\|_{\mathcal{V}^*}$ is indeed a norm and we consider the associated the dual norm, defined on \mathcal{V} , which we denote $\|\cdot\|_{\mathcal{V}}$. We can now consider the following generator of \mathcal{C}° : $\mathcal{X} = \mathcal{B} \cap \mathcal{C}^\circ$, where \mathcal{B} denotes the closed unit ball with respect to $\|\cdot\|_{\mathcal{V}}$. The following proposition shows that this choice of \mathcal{X} makes the average regret $\overline{\text{Reg}}_T$ bounded from above by $I_{\mathcal{X}}^*(\bar{r}_T)$.

Proposition 13 *Let $(a_t)_{t \geq 1}$ and $(\ell_t)_{t \geq 1}$ be sequences of actions chosen by the Decision Maker and the Environment respectively. Denote for all $t \geq 1$, $r_t = r(a_t, \ell_t)$ the corresponding payoffs. Then for all $T \geq 1$, the regret is bounded as*

$$\overline{\text{Reg}}_T = \left\| \frac{1}{T} \sum_{t=1}^T a_t \odot \ell_t \right\| - \min_{a \in \Delta_d} \left\| \frac{1}{T} \sum_{t=1}^T a \odot \ell_t \right\| \leq I_{\mathcal{B} \cap \mathcal{C}^\circ}^*(\bar{r}_T),$$

where \mathcal{B} denotes the closed unit ball associated with $\|\cdot\|_{\mathcal{V}}$.

Proof Let $T \geq 1$ and denote $y = \frac{1}{T} \sum_{t=1}^T a_t \odot \ell_t$ and $y' = \frac{1}{T} \sum_{t=1}^T \ell_t$. Let $(\tilde{y}, \tilde{y}') \in \mathcal{C}$ be any vector from the target set. Then, we can write

$$\begin{aligned}
 \overline{\text{Reg}}_T &= \|y\| - \min_{a \in \Delta_d} \|a \odot y'\| = \|y\| - \|\tilde{y}\| + \|\tilde{y}\| - \min_{a \in \Delta_d} \|a \odot y'\| \\
 &\quad + \min_{a \in \Delta_d} \|a \odot \tilde{y}'\| - \min_{a \in \Delta_d} \|a \odot \tilde{y}'\| \\
 &\leq \|y - y'\| + \min_{a \in \Delta_d} \|a \odot \tilde{y}'\| - \min_{a \in \Delta_d} \|a \odot y'\| \\
 &= \|y' - y\| + \max_{a \in \Delta_d} \min_{a' \in \Delta_d} \{\|a' \odot \tilde{y}'\| - \|a \odot y'\|\} \\
 &\leq \|y - y'\| + \max_{a \in \Delta_d} \|a \odot (\tilde{y}' - y')\| = \|(y, y') - (\tilde{y}, \tilde{y}')\|_{\mathcal{V}^*},
 \end{aligned}$$

where the first inequality follows from the reverse triangle inequality and the definition of \mathcal{C} and the third inequality from removing the minimum over $a' \in \Delta_d$ and using the reverse triangle inequality again. Then, taking the minimum over $(\tilde{y}, \tilde{y}') \in \mathcal{C}$ and applying Proposition 4 gives the result:

$$\overline{\text{Reg}}_T \leq \min_{(\tilde{y}, \tilde{y}') \in \mathcal{C}} \|(y, y') - (\tilde{y}, \tilde{y}')\|_{\mathcal{V}^*} = I_{\mathcal{B} \cap \mathcal{C}^\circ}^*(y, y') = I_{\mathcal{B} \cap \mathcal{C}^\circ} \left(\frac{1}{T} \sum_{t=1}^T r(a_t, \ell_t) \right).$$

■

4.4 An algorithm for ℓ_p global cost functions

We define and analyze an algorithm based on a carefully chosen regularizer which takes advantage of the properties of ℓ_p norms. The construction for general norms is given in Appendix F. We consider on $\mathcal{X} = \mathcal{B} \cap \mathcal{C}^\circ$ the following regularizer:

$$h(z, z') = \begin{cases} \frac{A}{2} \|z\|_2^2 + \frac{1}{2} \|z'\|_{q'}^2 & \text{if } (z, z') \in \mathcal{B} \cap \mathcal{C}^\circ, \\ +\infty & \text{otherwise.} \end{cases}$$

where $q' \in (1, 2]$ and $A > 0$ are to be chosen later. The algorithm associated with a positive sequence $(\eta_t)_{t \geq 1}$ and an oracle a from Remark 12 writes, for $t \geq 1$,

$$\begin{aligned}
 \text{compute } x_t &= \arg \max_{x \in \mathcal{X}} \left\{ \left\langle \eta_{t-1} \sum_{s=1}^{t-1} r_s, x \right\rangle - h(x) \right\} \\
 \text{compute } a_t &= \arg \min_{a \in \Delta_d} \sum_{i=1}^d \max(0, z_{ti} a_i + z'_{ti}), \quad \text{where } (z_t, z'_t) = x_t, \\
 \text{observe } r_t &:= r(a_t, b_t).
 \end{aligned}$$

Theorem 14 *Let $p \in (1, +\infty]$ and assume $\|\cdot\| = \|\cdot\|_p$. Then, the above algorithm with $A = \min\{d^{1-2/p}, 1\}$, $q' = 1 + (2 \log d - 1)^{-1}$ and coefficients*

$$\eta_t = \frac{1}{2\sqrt{t \max\{d^{2/p-1}, e(2 \log d - 1)\}}}, \quad t \geq 1,$$

guarantees, against any sequence $(\ell_t)_{t \geq 1}$ in $[0, 1]^d$ chosen by the Environment,

$$\forall T \geq 1, \quad \overline{\text{Reg}}_T \leq \frac{4}{\sqrt{T}} \max\left\{d^{1/p-1/2}, \sqrt{2e \log d}\right\}. \quad (2)$$

Remark 15 *In the special case $p = \infty$, the above bound recovers the best known bound of order $O(\sqrt{(\log d)/T})$ from [Rakhlin et al. \(2011\)](#). For $1 < p < +\infty$, we obtain, to the best of our knowledge, the first bounds with explicit dependence in d and p . Surprisingly, the same $O(\sqrt{(\log d)/T})$ bound with logarithmic dependence in the dimension d also holds for all $p \geq 2$. We were unable to find in the literature any lower bound for a given cost function³, and standard techniques from regret minimization, which involves a randomized Environment which cancels the influence of the Decision Maker on its own reward, do not seem to work at all, because of the particular form of the quantity to be minimized. Developing lower bound techniques for this kind of online learning problems appears to be an interesting and challenging research direction.*

Proof We aim at applying Theorem 10. Let us first establish an upper bound on the difference between the highest and lowest values of h . Note that $\max_{a \in \Delta_d} \|a \odot y'\|_p = \|y'\|_\infty$ for all $y' \in \mathbb{R}^d$. Indeed, by denoting e_1, \dots, e_d the canonical basis of \mathbb{R}^d , and using the fact that $\|\cdot\|_p \leq \|\cdot\|_1$,

$$\begin{aligned} \|y'\|_\infty &= \max_{1 \leq i \leq d} |y'_i| = \max_{a \in \{e_1, \dots, e_d\}} \|a \odot y'\|_p \leq \max_{a \in \Delta_d} \|a \odot y'\|_p \\ &\leq \max_{a \in \Delta_d} \|a \odot y'\|_1 = \max_{a \in \Delta_d} \sum_{i=1}^d a_i |y'_i| = \|y'\|_\infty. \end{aligned}$$

Therefore, $\|(y, y')\|_{\mathcal{V}^*} = \|y\|_p + \|y'\|_\infty$ for all $(y, y') \in \mathcal{V}^*$. Using a standard argument, we can prove that its dual norm writes

$$\|(z, z')\|_{\mathcal{V}} = \max\left\{\|z\|_q, \|z'\|_1\right\}, \quad (z, z') \in \mathcal{V},$$

where $q = (1 - 1/p)^{-1}$. Therefore, $\mathcal{B} = \left\{(z, z') \in \mathcal{V}, \|z\|_q \leq 1 \text{ and } \|z'\|_1 \leq 1\right\}$. Besides, because $0 \in \mathcal{X}$, it holds that $\min_{\mathcal{X}} h = 0$. Therefore, using the standard inequality between

³. [Even-Dar et al. \(2009\)](#) gives a lower bound, but is of a different kind, as the cost function depends on the time horizon.

ℓ_p norms that can be written $\|\cdot\|_{q'} \leq d^{\max(1/q'-1/q,0)} \|\cdot\|_q$,

$$\begin{aligned} \max_{\mathcal{X}} h - \min_{\mathcal{X}} h &\leq \max_{(z,z') \in \mathcal{B}} h(z, z') = \max_{\substack{\|z\|_q \leq 1 \\ \|z'\|_1 \leq 1}} \left\{ \frac{A}{2} \|z\|_2^2 + \frac{1}{2} \|z'\|_{q'}^2 \right\} \\ &\leq \max_{\substack{\|z\|_q \leq 1 \\ \|z'\|_1 \leq 1}} \left\{ \frac{A}{2} d^{\max(1-2/q,0)} \|z\|_q^2 + \frac{1}{2} \|z'\|_1^2 \right\} \\ &= \frac{1}{2} (A d^{\max(1-2/q,0)} + 1) = \frac{1}{2} (A d^{\max(2/p-1,0)} + 1). \end{aligned} \quad (3)$$

Let us introduce the following norm on the payoff space \mathcal{V}^* , which is different from the norm $\|\cdot\|_{\mathcal{V}^*}$ involved in Proposition 13:

$$\|(y, y')\|_{(\mathcal{V}^*)} = \|y\|_1 + \|y'\|_\infty.$$

We can see that the vector-valued payoffs are bounded by 2 with respect to this norm. Indeed, for all $a \in \Delta_d$ and $\ell \in [0, 1]^d$,

$$\|r(a, \ell)\|_{(\mathcal{V}^*)} = \|a \odot \ell\|_1 + \|\ell\|_\infty \leq 2.$$

Denote $\|\cdot\|_{(\mathcal{V})}$ the dual norm of $\|\cdot\|_{(\mathcal{V}^*)}$, which has the following expression: $\|(z, z')\|_{(\mathcal{V})} = \max(\|z\|_\infty, \|z'\|_1)$ for all $(z, z') \in \mathcal{V}$.

Let us now prove for regularizer h a strong convexity property with respect to $\|\cdot\|_{(\mathcal{V})}$. It can be practical to write h as

$$h(z, z') = h_1(z) + h_2(z') + I_{\mathcal{X}}(z, z'), \quad (z, z') \in \mathcal{V},$$

where $h_1(z) = \frac{A}{2} \|z\|_2^2$ and $h_2(z) = \frac{1}{2} \|z'\|_{q'}^2$. We note that according to Proposition 24, h_1 is A -strongly convex with respect to $\|\cdot\|_2$ and h_2 is $(q' - 1)d^{2(1/q'-1)}$ -strongly convex with respect to $\|\cdot\|_1$. For all $(z, z'), (\tilde{z}, \tilde{z}') \in \mathcal{V}$, and $\lambda \in [0, 1]$, denote $z_\lambda = \lambda z + (1 - \lambda)\tilde{z}$ and $z'_\lambda = \lambda z' + (1 - \lambda)\tilde{z}'$. Then, using the strong convexity properties of h_1 and h_2 , and the fact that $\|\cdot\|_2 \geq \|\cdot\|_\infty$,

$$\begin{aligned} \lambda h(z, z') + (1 - \lambda)h(\tilde{z}, \tilde{z}') &\geq \lambda(h_1(z) + h_2(z')) + (1 - \lambda)(h_1(\tilde{z}) + h_2(\tilde{z}')) \\ &= \lambda h_1(z) + (1 - \lambda)h_1(\tilde{z}) + \lambda h_2(z') + (1 - \lambda)h_2(\tilde{z}') \\ &\geq h_1(z_\lambda) + \frac{A\lambda(1 - \lambda)}{2} \|\tilde{z} - z\|_2^2 + h_2(z'_\lambda) \\ &\quad + \frac{(q' - 1)d^{2(1/q'-1)}\lambda(1 - \lambda)}{2} \|\tilde{z}' - z'\|_1^2 \\ &\geq h(z_\lambda, z'_\lambda) \\ &\quad + \min \left\{ A, (q' - 1)d^{2(1/q'-1)} \right\} \frac{\lambda(1 - \lambda)}{2} \|(\tilde{z}, \tilde{z}') - (z, z')\|_{(\mathcal{V})}^2. \end{aligned}$$

Therefore, h is $\min \left\{ A, (q' - 1)d^{2(1/q'-1)} \right\}$ -strongly convex with respect to $\|\cdot\|_{(\mathcal{V})}$.

Applying Theorem 10 with

$$M = 2, \quad \Delta = \frac{1}{2}(A d^{\max(2/p-1,0)} + 1), \quad \text{and} \quad K = \min \left\{ A, (q' - 1)d^{2/q'-2} \right\},$$

together with Proposition 13 gives

$$\overline{\text{Reg}}_T \leq \frac{4}{\sqrt{2}} \sqrt{\frac{A d^{\max(2/p-1,0)} + 1}{T \min \left\{ A, (q' - 1)d^{2/q'-2} \right\}}} = \frac{4}{\sqrt{T}} \max \left\{ d^{1/p-1/2}, \sqrt{e(2 \log d - 1)} \right\},$$

where the equality follows from the choice $A = \min \{d^{1-2/p}, 1\}$ and $q' = 1 + (2 \log d - 1)^{-1}$. Hence the result. \blacksquare

Acknowledgments

The author is grateful to Rida Laraki, Vianney Perchet, Sylvain Sorin and Gilles Stoltz for inspiring discussions, insightful advice and careful proofreading. The paper also benefited from the numerous suggestions of improvements from the two anonymous referees. This work was supported by a public grant as part of the *Investissement d'avenir* project, reference ANR-11-LABX-0056-LMH, LabEx LMHA.

References

- J. Abernethy, P. L. Bartlett, and E. Hazan. Blackwell approachability and low-regret learning are equivalent. In *JMLR: Workshop and Conference Proceedings (COLT)*, volume 19, pages 27–46, 2011.
- Y. Azar, B. Kalyanasundaram, S. Plotkin, K. R. Pruhs, and O. Waarts. Online load balancing of temporary tasks. In *Workshop on algorithms and data structures*, pages 119–130. Springer, 1993.
- Y. Azar, U. Felge, M. Feldman, and M. Tennenholtz. Sequential decision making with vector outcomes. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 195–206, 2014.
- A. Bernstein and N. Shimkin. Response-based approachability with applications to generalized no-regret problems. *The Journal of Machine Learning Research*, 16(1):747–773, 2015.
- D. Blackwell. Controlled random walks. In *Proceedings of the International Congress of Mathematicians*, volume 3, pages 336–338, 1954.
- D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
- A. Blum and Y. Mansour. From external to internal regret. In *Learning Theory*, pages 621–636. Springer, 2005.

- A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.
- J. M. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer, 2010.
- S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- S. Bubeck. *Introduction to Online Optimization: Lecture Notes*. Princeton University, 2011.
- N. Cesa-Bianchi. Analysis of two gradient-based algorithms for on-line regression. In *Proceedings of the Tenth Annual Conference on Computational Learning Theory (COLT)*, pages 163–170. ACM, 1997.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- A. P. Dawid. The well-calibrated bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982.
- E. Even-Dar, R. Kleinberg, S. Mannor, and Y. Mansour. Online learning for global cost functions. In *COLT*, 2009.
- G. Farina, C. Kroer, and T. Sandholm. Faster game solving via predictive Blackwell approachability: Connecting regret matching and mirror descent. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5363–5371, May 2021.
- D. P. Foster and R. V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1):40–55, 1997.
- D. P. Foster and R. V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.
- D. Fudenberg and D. K. Levine. Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19(5):1065–1089, 1995.
- D. Fudenberg and D. K. Levine. Conditional universal consistency. *Games and Economic Behavior*, 29(1):104–130, 1999.
- C. Gentile and M. K. Warmuth. Linear hinge loss and average margin. In *Advances in Neural Information Processing Systems (NIPS)*, volume 11, pages 225–231, 1998.
- G. J. Gordon. No-regret algorithms for online convex programs. In *Advances in Neural Information Processing Systems*, pages 489–496, 2007.
- A. J. Grove, N. Littlestone, and D. Schuurmans. General convergence results for linear discriminant updates. *Machine Learning*, 43(3):173–210, 2001.
- J. Hannan. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, 3(97–139):2, 1957.
- S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150, 2000.

- S. Hart and A. Mas-Colell. A general class of adaptive strategies. *Journal of Economic Theory*, 98(1):26–54, 2001.
- E. Hazan, S. Kale, and M. K. Warmuth. Learning rotations with little regret. In *Proceedings of the 23rd Conference on Learning Theory (COLT)*, pages 144–154, 2010.
- D. P. Helmbold and M. K. Warmuth. Learning permutations with exponential weights. *The Journal of Machine Learning Research*, 10:1705–1736, 2009.
- S. M. Kakade, S. Shalev-Shwartz, and A. Tewari. Regularization techniques for learning with matrices. *The Journal of Machine Learning Research*, 13(1):1865–1890, 2012.
- A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- J. Kivinen and M. K. Warmuth. Relative loss bounds for multidimensional regression problems. *Machine Learning*, 45(3):301–329, 2001.
- E. Kohlberg. Optimal strategies in repeated games with incomplete information. *International Journal of Game Theory*, 4(1):7–24, 1975.
- W. M. Koolen, M. K. Warmuth, and J. Kivinen. Hedging structured concepts. In *Proceedings of the 23rd Conference on Learning Theory (COLT)*, pages 93–105, 2010.
- Y. Liu, K. Hatano, and E. Takimoto. Improved algorithms for online load balancing. In *SOFSEM 2021: Theory and Practice of Computer Science*, pages 203–217. Springer International Publishing, 2021.
- S. Mannor and N. Shimkin. Regret minimization in repeated matrix games with variable stage duration. *Games and Economic Behavior*, 63(1):227–258, 2008.
- S. Mannor, V. Perchet, and G. Stoltz. Approachability in unknown games: Online learning meets multi-objective optimization. In *Conference on Learning Theory*, pages 339–355, 2014.
- M. Molinaro. Online and random-order load balancing simultaneously. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1638–1650. SIAM, 2017.
- J.-J. Moreau. Décomposition orthogonale d’un espace hilbertien selon deux cônes mutuellement polaires. *Comptes rendus de l’Académie des Sciences*, 255:238–240, 1962.
- Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical programming*, 120(1):221–259, 2009.
- V. Perchet. Approachability, regret and calibration: Implications and equivalences. *Journal of Dynamics and Games*, 1(2):181–254, 2014.
- V. Perchet. Exponential weight approachability, applications to calibration and regret minimization. *Dynamic Games and Applications*, 5(1):136–153, 2015.

- A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Beyond regret. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 559–594, 2011.
- R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- S. Shalev-Shwartz. *Online learning: Theory, Algorithms, and Applications*. PhD thesis, The Hebrew University of Jerusalem, 2007.
- S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- N. Shimkin. An online convex optimization approach to Blackwell’s approachability. *The Journal of Machine Learning Research*, 17(1):4434–4456, 2016.
- G. Stoltz and G. Lugosi. Internal regret in on-line portfolio selection. *Machine Learning*, 59(1-2):125–159, 2005.
- E. Takimoto and M. K. Warmuth. Path kernels and multiplicative updates. *The Journal of Machine Learning Research*, 4:773–818, 2003.
- O. Tammelin, N. Burch, M. Johanson, and M. Bowling. Solving heads-up limit texas hold’em. In *Twenty-fourth international joint conference on artificial intelligence*, 2015.
- M. K. Warmuth and D. Kuzmin. Randomized online PCA algorithms with regret bounds that are logarithmic in the dimension. *Journal of Machine Learning Research*, 9(10): 2287–2320, 2008.
- M. Zinkevich, M. Johanson, M. Bowling, and C. Piccione. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20: 1729–1736, 2007.




Figure 1: The polar cone of a set \mathcal{A} and the bipolar

Appendix A. Definitions and properties about closed convex cones

We recall the definitions of a closed convex cone, of the polar cone, and gather a few properties. \mathcal{W} will be a finite-dimensional vector space and \mathcal{W}^* its dual.

Definition 16 *A nonempty subset \mathcal{C} of \mathcal{W} is a closed convex cone if it is closed and if for all $y, y' \in \mathcal{C}$ and $\lambda \in \mathbb{R}_+$, we have $y + y' \in \mathcal{C}$ and $\lambda y \in \mathcal{C}$.*

Definition 17 *Let \mathcal{A} be a subset of \mathcal{W} . The polar cone of \mathcal{A} is a subset of the dual space \mathcal{W}^* defined by*

$$\mathcal{A}^\circ = \{x \in \mathcal{W}^*, \forall y \in \mathcal{A}, \langle y, x \rangle \leq 0\}.$$

The following proposition is an immediate consequence of the bipolar theorem—see e.g. Theorem 3.3.14 in [Borwein and Lewis \(2010\)](#).

Proposition 18 *Let \mathcal{A} be a subset of \mathcal{W} .*

- (i) $\mathcal{A}^{\circ\circ}$ is the smallest closed convex cone containing \mathcal{A} .
- (ii) If \mathcal{A} is closed and convex, then $\mathcal{A}^{\circ\circ} = \mathbb{R}_+\mathcal{A}$.
- (iii) If \mathcal{A} is a closed convex cone, then $\mathcal{A}^{\circ\circ} = \mathcal{A}$.

The following statement is a simpler version of Moreau's decomposition theorem ([Moreau, 1962](#)).

Proposition 19 *Assume that \mathcal{W} is an Euclidean space. We identify \mathcal{W} and its dual space \mathcal{W}^* . Let \mathcal{C} be a closed convex cone in \mathcal{W} , and $y \in \mathcal{W}$. Then, $y - \text{proj}_{\mathcal{C}} y = \text{proj}_{\mathcal{C}^\circ} y$, where proj denotes the Euclidean projection. In particular, $y - \text{proj}_{\mathcal{C}} y$ belongs to \mathcal{C}° .*

A.1 Proof of Proposition 2

(i) is easy. (ii) holds because $\mathcal{B} \cap \mathcal{C}$ is indeed nonempty, convex as the intersection of two convex sets, and for any point $x \in \mathcal{C} \setminus \{0\}$, $x/\|x\|$ belongs to $\mathcal{B} \cap \mathcal{C}$, so that $\mathbb{R}_+(\mathcal{B} \cap \mathcal{C}) = \mathcal{C}$. (iii) is a consequence of Proposition 18.




Figure 2: Illustration of Proposition 19

Appendix B. Properties of regularizers

Proposition 20 *Let h be a regularizer on \mathcal{X} . Its Legendre–Fenchel transform, defined by*

$$h^*(y) = \sup_{x \in \mathcal{V}} \{\langle y, x \rangle - h(x)\}, \quad y \in \mathcal{V}^*,$$

satisfies the following properties.

- (i) $\text{dom } h^* = \mathcal{V}^*$;
- (ii) h^* is differentiable on \mathcal{V}^* ;
- (iii) For all $y \in \mathcal{V}^*$, $\nabla h^*(y) = \arg \max_{x \in \mathcal{X}} \{\langle y, x \rangle - h(x)\}$. In particular, ∇h^* takes values in \mathcal{X} .

Proof (i) Let $w \in \mathcal{V}^*$. The function $x \mapsto \langle w, x \rangle - h(x)$ equals $-\infty$ outside of \mathcal{X} , and is upper semicontinuous on \mathcal{X} which is compact. It thus has a maximum and $h^*(w) < +\infty$.

(ii,iii) Moreover, this maximum is attained at a unique point because h is strictly convex. Besides, for $x \in \mathcal{V}$ and $w \in \mathcal{V}^*$

$$x \in \partial h^*(w) \iff w \in \partial h(x) \iff x \in \arg \max_{x' \in \mathcal{X}} \{\langle w, x' \rangle - h(x')\},$$

in other words, $\partial h^*(w) = \arg \max_{x' \in \mathcal{X}} \{\langle w, x' \rangle - h(x')\}$. This argmax is a singleton as we noticed. It means that h^* is differentiable. ■

Recall that Δ_d denotes the unit simplex of \mathbb{R}^d : $\Delta_d = \left\{ x \in \mathbb{R}_+^d \mid \sum_{i=1}^d x_i = 1 \right\}$.

Definition 21 (Entropic regularizer) *The entropic regularizer $h_{\text{ent}} : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is defined as*

$$h_{\text{ent}}(x) = \begin{cases} \sum_{i=1}^d x_i \log x_i & \text{if } x \in \Delta_d \\ +\infty & \text{otherwise,} \end{cases}$$

where $x_i \log x_i = 0$ when $x_i = 0$.

Proposition 22 (i) h_{ent} is a regularizer on Δ_d ;

$$(ii) \nabla h_{\text{ent}}^*(y) = \left(\frac{\exp y_i}{\sum_{j=1}^d \exp y_j} \right)_{1 \leq i \leq d}, \text{ for all } y \in \mathbb{R}^d;$$

$$(iii) \max_{x \in \Delta_d} h_{\text{ent}}(x) - \min_{x \in \Delta_d} h_{\text{ent}}(x) = \log d;$$

(iv) h_{ent} is 1-strongly convex with respect to $\|\cdot\|_1$.

Proof (i) is immediate, and (ii) is classic—see e.g. (Boyd and Vandenberghe, 2004, Example 2.25).

(iii) h_{ent} being convex, its maximum on Δ_d is attained at one of the extreme points. At each extreme point, the value of h_{ent} is zero. Therefore, $\max_{\Delta_d} h_{\text{ent}} = 0$. As for the minimum, h_{ent} being convex and symmetric with respect to the components x_i , its minimum is attained at the centroid $(1/d, \dots, 1/d)$ of the simplex Δ_d , where its value is $-\log d$. Therefore, $\min_{\Delta_d} h_{\text{ent}} = -\log d$ and $\max_{\Delta_d} h_{\text{ent}} - \min_{\Delta_d} h_{\text{ent}} = \log d$.

(iv) Consider $F : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ defined by

$$F(x) = \begin{cases} \sum_{i=1}^d (x_i \log x_i - x_i) + 1 & \text{if } x \in \mathbb{R}_+^d \\ +\infty & \text{otherwise.} \end{cases}$$

Let us prove that F is 1-strongly convex with respect to $\|\cdot\|_1$. By definition, the domain of F is \mathbb{R}_+^d . It is differentiable on the interior of the domain $(\mathbb{R}_+^*)^d$ and $\nabla F(x) = (\log x_i)_{1 \leq i \leq d}$ for $x \in (\mathbb{R}_+^*)^d$. Therefore, the norm of $\nabla F(x)$ goes to $+\infty$ when x converges to a boundary point of \mathbb{R}_+^d . (Rockafellar, 1970, Theorem 26.1) then assures that the subdifferential $\partial F(x)$ is empty as soon as $x \notin (\mathbb{R}_+^*)^d$. Therefore, the characterization of strong convexity from (Shalev-Shwartz, 2007, Lemma 14), which we aim at proving, can be written

$$\langle \nabla F(x') - \nabla F(x), x' - x \rangle \geq \|x' - x\|_1^2, \quad x, x' \in (\mathbb{R}_+^*)^d. \quad (4)$$

Let $x, x' \in (\mathbb{R}_+^*)^d$.

$$\langle \nabla F(x') - \nabla F(x), x' - x \rangle = \sum_{i=1}^d \log \frac{x'_i}{x_i} (x'_i - x_i).$$

A simple study of function shows that $(s-1) \log s - 2(s-1)^2/(s+1) \geq 0$ for $s \geq 0$. Applied with $s = x'_i/x_i$, this gives

$$\sum_{i=1}^d \log \frac{x'_i}{x_i} (x'_i - x_i) \geq \|x' - x\|_1^2,$$

and (4) is proved. F is therefore 1-strongly convex with respect to $\|\cdot\|_1$ and so is h_{ent} thanks to Lemma 25. ■

Definition 23 (ℓ_p regularizer) For $p \in (1, 2]$ and a nonempty convex compact subset \mathcal{X} of \mathbb{R}^d , the associated ℓ_p regularizer is defined as

$$h_p(x) = \begin{cases} \frac{1}{2} \|x\|_p^2 & \text{if } x \in \mathcal{X} \\ +\infty & \text{otherwise.} \end{cases}$$

Proposition 24 Let $p \in (1, 2]$.

- (i) h_p is a regularizer on \mathcal{X} ;
- (ii) h_p is $(p-1)d^{2(1/p-1)}$ -strongly convex with respect to $\|\cdot\|_1$;
- (iii) h_2 is 1-strongly convex with respect to $\|\cdot\|_2$;
- (iv) $\nabla h_2^*(y) = \text{proj}_{\mathcal{X}}(y)$ for all $y \in \mathbb{R}^d$ where $\text{proj}_{\mathcal{X}}$ denotes the Euclidean projection onto \mathcal{X} .

Proof (i) Since $p \geq 1$, $\|\cdot\|_p$ is a norm and is therefore convex. h_p then clearly is a regularizer on \mathcal{X} . (ii,iii) We consider the function $F(x) = \frac{1}{2} \|x\|_p^2$ defined on \mathbb{R}^d which is $(p-1)$ -strongly convex with respect to $\|\cdot\|_p$ —see e.g. Bubeck (2011) or (Kakade et al., 2012, Corollary 10). Then, so is h_p thanks to Lemma 25. Substituting $p = 2$ gives (iii). The strong convexity with respect to $\|\cdot\|_1$ follows from the standard comparison $\|\cdot\|_p \geq d^{1/q-1} \|\cdot\|_1$ in \mathbb{R}^d . (iv) For all $y \in \mathbb{R}^d$, using property (iii) from Proposition 20,

$$\begin{aligned} \nabla h_2^*(y) &= \arg \max_{x \in \mathcal{X}} \left\{ \langle y, x \rangle - \frac{1}{2} \|x\|_2^2 \right\} = \arg \min_{x \in \mathcal{X}} \left\{ \frac{1}{2} \|x\|_2^2 - \langle y, x \rangle + \frac{1}{2} \|y\|_2^2 \right\} \\ &= \arg \min_{x \in \mathcal{X}} \|y - x\|_2^2 = \text{proj}_{\mathcal{X}}(y). \end{aligned}$$

■

Lemma 25 Let $\|\cdot\|$ a norm on \mathcal{V} , $K > 0$ and $h, F : \mathcal{V} \rightarrow \mathbb{R} \cup \{+\infty\}$ two convex functions such that for all $x \in \mathcal{V}$,

$$h(x) = F(x) \quad \text{or} \quad h(x) = +\infty.$$

Then, if F is K -strongly convex with respect to $\|\cdot\|$, so is h .

Proof Note that for all $x \in \mathcal{V}$, $F(x) \leq h(x)$. Let us prove that h satisfies the condition from Definition 8. Let $x, x' \in \mathcal{V}$, $\lambda \in [0, 1]$ and denote $x'' = \lambda x + (1-\lambda)x'$. Let us first assume that $h(x'') = +\infty$. By convexity of h , either $h(x)$ or $h(x')$ is equal to $+\infty$, and the right-hand side of (1) is equal to $+\infty$. Inequality (1) therefore holds. If $h(x'')$ is finite,

$$\begin{aligned} h(x'') &= F(x'') \leq \lambda F(x) + (1-\lambda)F(x') - \frac{K\lambda(1-\lambda)}{2} \|x' - x\|^2 \\ &\leq \lambda h(x) + (1-\lambda)h(x') - \frac{K\lambda(1-\lambda)}{2} \|x' - x\|^2, \end{aligned}$$

and (1) is proved. ■

Appendix C. Various postponed proofs

C.1 Proof of Proposition 4

Let $y \in \mathcal{V}^*$. Using the definition of the dual norm and Sion's minimax theorem,

$$\inf_{y' \in \mathcal{C}} \|y' - y\|_* = \inf_{y' \in \mathcal{C}} \sup_{x \in \mathcal{B}} \langle y - y', x \rangle = \sup_{x \in \mathcal{B}} \inf_{y' \in \mathcal{C}} \{ \langle y, x \rangle - \langle y', x \rangle \}.$$

Suppose x does not belong to \mathcal{C}° . Then, there exists $y'_0 \in \mathcal{C}$ such that $\langle y'_0, x \rangle > 0$. \mathcal{C} being closed by multiplication by \mathbb{R}_+ , the quantity $\langle y', x \rangle$ (with $y' \in \mathcal{C}$) can be made arbitrarily large by selecting $y' = \lambda y'_0$ and letting $\lambda \rightarrow +\infty$, and thus the above infimum is equal to $-\infty$. Therefore, we can restrict the above supremum to $\mathcal{B} \cap \mathcal{C}^\circ$. We thus have

$$\inf_{y' \in \mathcal{C}} \|y' - y\|_* = \sup_{x \in \mathcal{B} \cap \mathcal{C}^\circ} \left\{ \langle y, x \rangle - \sup_{y' \in \mathcal{C}} \langle y', x \rangle \right\}.$$

The above embedded supremum is zero because for $x \in \mathcal{B} \cap \mathcal{C}^\circ$ and $y' \in \mathcal{C}$ we obviously have $\langle y', x \rangle \leq 0$, and 0 is attained with $y' = 0$. Finally,

$$\inf_{y' \in \mathcal{C}} \|y' - y\|_* = \sup_{x \in \mathcal{B} \cap \mathcal{C}^\circ} \langle y, x \rangle = I_{\mathcal{B} \cap \mathcal{C}^\circ}^*(y).$$

C.2 Proof of Proposition 6

Blackwell's condition can be written

$$\max_{x \in \mathcal{C}^\circ} \min_{a \in \mathcal{A}} \max_{b \in \mathcal{B}} \langle r(a, b), x \rangle \leq 0.$$

Since the above dot product is affine in each of the variables a , b and x , by applying Sion's minimax theorem twice, the above is equivalent to

$$\max_{b \in \mathcal{B}} \min_{a \in \mathcal{A}} \max_{x \in \mathcal{C}^\circ} \langle r(a, b), x \rangle \leq 0,$$

which is exactly the dual condition.

C.3 Proof of Lemma 9

Assume that the sequence of parameters $(\eta_t)_{t \geq 1}$ is nonincreasing. Denote $Y_t = \sum_{s=1}^t r_s$ for $t \geq 1$ and $\eta_0 = \eta_1$. Let $x \in \mathcal{X}$. Using Fenchel's inequality, we write

$$\begin{aligned} \langle Y_T, x \rangle &= \frac{\langle \eta_T Y_T, x \rangle}{\eta_T} \leq \frac{h^*(\eta_T Y_T)}{\eta_T} + \frac{h(x)}{\eta_T} \\ &\leq \frac{h^*(0)}{\eta_0} + \sum_{t=1}^T \left(\frac{h^*(\eta_t Y_t)}{\eta_t} - \frac{h^*(\eta_{t-1} Y_{t-1})}{\eta_{t-1}} \right) + \frac{\max_{x \in \mathcal{X}} h(x)}{\eta_T}. \end{aligned} \tag{5}$$

Let us bound $h^*(\eta_t Y_t)/\eta_t$ from above. For all $x \in \mathcal{X}$ we have

$$\frac{\langle \eta_t Y_t, x \rangle - h(x)}{\eta_t} = \frac{\langle \eta_{t-1} Y_t, x \rangle - h(x)}{\eta_{t-1}} - h(x) \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right).$$

The maximum over $x \in \mathcal{X}$ of the above left-hand side gives $h^*(\eta_t Y_t)/\eta_t$. As for the right-hand side, let us take the maximum over $x \in \mathcal{X}$ for each of the two terms separately. This gives

$$\begin{aligned} \frac{h^*(\eta_t Y_t)}{\eta_t} &\leq \max_{x \in \mathcal{X}} \left\{ \frac{\langle \eta_{t-1} Y_t, x \rangle - h(x)}{\eta_{t-1}} \right\} + \max_{x \in \mathcal{X}} \left\{ -h(x) \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \right\} \\ &= \frac{h^*(\eta_{t-1} Y_t)}{\eta_{t-1}} + \left(\min_{x \in \mathcal{X}} h(x) \right) \left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right), \end{aligned}$$

where we used the fact that the sequence $(\eta_t)_{t \geq 0}$ is nonincreasing. Injecting this inequality in (5), we get

$$\begin{aligned} \langle Y_T, x \rangle &\leq \frac{h^*(0)}{\eta_0} + \sum_{t=1}^T \frac{h^*(\eta_{t-1} Y_t) - h^*(\eta_{t-1} Y_{t-1})}{\eta_{t-1}} \\ &\quad + \left(\min_{x \in \mathcal{X}} h(x) \right) \sum_{t=1}^T \left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) + \frac{\max_{x \in \mathcal{X}} h(x)}{\eta_T}. \end{aligned}$$

We now make the quantity

$$D_{h^*}(\eta_{t-1} Y_t, \eta_{t-1} Y_{t-1}) := h^*(\eta_{t-1} Y_t) - h^*(\eta_{t-1} Y_{t-1}) - \langle \nabla h^*(\eta_{t-1} Y_{t-1}), \eta_{t-1} Y_t - \eta_{t-1} Y_{t-1} \rangle$$

(called a Bregman divergence) appear in the first above sum by subtracting

$$\frac{\langle \eta_{t-1} Y_t - \eta_{t-1} Y_{t-1}, \nabla h^*(\eta_{t-1} Y_{t-1}) \rangle}{\eta_{t-1}} = \langle r_t, x_t \rangle.$$

Therefore,

$$\begin{aligned} \langle Y_T, x \rangle &\leq \frac{h^*(0)}{\eta_0} + \sum_{t=1}^T \frac{D_{h^*}(\eta_{t-1} Y_t, \eta_{t-1} Y_{t-1})}{\eta_{t-1}} + \sum_{t=1}^T \langle r_t, x_t \rangle \\ &\quad - \frac{\min_{x \in \mathcal{X}} h(x)}{\eta_T} + \frac{\min_{x \in \mathcal{X}} h(x)}{\eta_0} + \frac{\max_{x \in \mathcal{X}} h(x)}{\eta_T}. \end{aligned}$$

Since $h^*(0) = -\min_{x \in \mathcal{X}} h(x)$, we get

$$\begin{aligned} \text{Reg}_T &= \max_{x \in \mathcal{X}} \langle Y_T, x \rangle - \sum_{t=1}^T \langle r_t, x_t \rangle \\ &\leq \frac{\max_{\mathcal{X}} h - \min_{x \in \mathcal{X}} h(x)}{\eta_T} + \sum_{t=1}^T \frac{D_{h^*}(\eta_{t-1} Y_t, \eta_{t-1} Y_{t-1})}{\eta_{t-1}} \\ &\leq \frac{\Delta}{\eta_T} + \sum_{t=1}^T \frac{D_{h^*}(\eta_{t-1} Y_t, \eta_{t-1} Y_{t-1})}{\eta_{t-1}}. \end{aligned}$$

The strong convexity of the regularizer h let us bound the above Bregman divergences as follows—see e.g. (Shalev-Shwartz, 2007, Lemma 13):

$$D_{h^*}(\eta_{t-1} Y_t, \eta_{t-1} Y_{t-1}) \leq \frac{1}{2K} \|\eta_{t-1} Y_t - \eta_{t-1} Y_{t-1}\|_*^2 = \frac{\eta_{t-1}^2}{2K} \|r_t\|_*^2, \quad t \geq 1.$$

Then, set $\eta = \sqrt{\Delta/M^2}$ so that $\eta_t = \eta t^{-1/2}$ for $t \geq 1$, which is indeed a nonincreasing sequence. The regret bound then becomes

$$\frac{\Delta\sqrt{T}}{\eta} + \frac{M^2}{2K} \sum_{t=1}^T \eta_{t-1}.$$

We bound the above sum as follows. Since $\eta_0 = \eta_1 = \eta$,

$$\begin{aligned} \sum_{t=1}^T \eta_{t-1} &= \eta \left(2 + \sum_{t=2}^{T-1} \frac{1}{\sqrt{t}} \right) \leq \eta \left(\int_0^1 \frac{1}{\sqrt{s}} ds + \int_1^{T-1} \frac{1}{\sqrt{s}} ds \right) \\ &= \eta \int_0^{T-1} \frac{1}{\sqrt{s}} ds = 2\eta\sqrt{T-1} \leq 2\eta\sqrt{T}. \end{aligned}$$

Injecting the expression of η and simplifying gives the result:

$$\text{Reg}_T \leq 2M\sqrt{\frac{T\Delta}{K}}.$$

Appendix D. Blackwell's algorithm

We recall the definition of Blackwell's algorithm (Blackwell, 1956) and show that it belongs to the family of FTRL algorithms defined in Section 3.2. In the related work by Shimkin (2016), it is demonstrated that Blackwell's algorithm can also be interpreted as a Follow the Leader algorithm, as well as a FTRL algorithm, in the context of online convex optimization algorithms converted into algorithms for the approachability of bounded convex target sets.

We consider $\mathcal{V} = \mathcal{V}^* = \mathbb{R}^d$ equipped with its Euclidean structure. Let $\mathcal{C} \subset \mathbb{R}^d$ be a closed convex cone which we assume to be a B-set for the game $(\mathcal{A}, \mathcal{B}, r)$ and $a : \mathcal{C}^\circ \rightarrow \mathcal{X}$ a $(\mathcal{A}, \mathcal{B}, r, \mathcal{C})$ -oracle. It follows from Definition 5 that it is always possible to choose an oracle a that satisfies

$$x = \lambda x' \text{ for some } \lambda > 0 \implies a(x) = a(x'), \quad x, x' \in \mathcal{C}^\circ. \quad (6)$$

We assume in this section that oracle a satisfies this property.

Blackwell's algorithm (Blackwell, 1954) is defined by

$$a_t = a \left(\bar{r}_{t-1} - \text{proj}_{\mathcal{C}} \bar{r}_{t-1} \right), \quad t \geq 1,$$

where $\text{proj}_{\mathcal{C}}$ denotes the Euclidean projection onto \mathcal{C} . It can be rewritten, using Proposition 19, as

$$a_t = a \left(\text{proj}_{\mathcal{C}^\circ} \bar{r}_{t-1} \right), \quad t \geq 1.$$

Theorem 26 *Let $\mathcal{X} = \mathcal{C}^\circ \cap \mathcal{B}$ where \mathcal{B} denotes the closed Euclidean ball, and h_2 the Euclidean regularizer on \mathcal{X} . Blackwell's algorithm and the FTRL algorithm associated with h_2 and any sequence of positive parameters $(\eta_t)_{t \geq 1}$ coincide. In other words,*

$$a \left(\bar{r}_{t-1} - \text{proj}_{\mathcal{C}} \bar{r}_{t-1} \right) = a \left(\nabla h_2^* \left(\eta_{t-1} \sum_{s=1}^{t-1} r_s \right) \right), \quad t \geq 1.$$




Figure 3: In the case where $\|\text{proj}_{\mathcal{C}^\circ} w\|_2 \leq 1$, we have $\text{proj}_{\mathcal{C}^\circ} w = \text{proj}_{\mathcal{C}^\circ \cap \mathcal{B}} w$.

Proof Recall that the Euclidean projection $\text{proj}_{\mathcal{E}} w$ of a point w on a closed convex set \mathcal{E} is the only point in \mathcal{E} satisfying

$$\forall w' \in \mathcal{E}, \quad \left\langle w - \text{proj}_{\mathcal{E}} w, w' - \text{proj}_{\mathcal{E}} w \right\rangle \leq 0. \quad (7)$$

This characterization will be needed later.

Remember from Proposition 24 that $\nabla h_2^* = \text{proj}_{\mathcal{C}^\circ \cap \mathcal{B}}$. Since oracle a satisfies property (6), it is enough to prove that for all $u \in \mathbb{R}^d$ and $\mu > 0$,

$$\text{proj}_{\mathcal{C}^\circ} u \in \mathbb{R}_+^* \text{proj}_{\mathcal{C}^\circ \cap \mathcal{B}}(\mu u).$$

Besides, \mathcal{C}° being a closed convex cone, $\text{proj}_{\mathcal{C}^\circ}(\mu u) = \mu \text{proj}_{\mathcal{C}^\circ} u$. It is therefore equivalent to prove that for all $w \in \mathbb{R}^d$,

$$\text{proj}_{\mathcal{C}^\circ} w \in \mathbb{R}_+^* \text{proj}_{\mathcal{C}^\circ \cap \mathcal{B}} w. \quad (8)$$

Let $w \in \mathbb{R}^d$. If $\|\text{proj}_{\mathcal{C}^\circ} w\|_2 \leq 1$, then obviously $\text{proj}_{\mathcal{C}^\circ} w = \text{proj}_{\mathcal{C}^\circ \cap \mathcal{B}} w$ as shown in Figure 3 and (8) is true. We now assume that $\|\text{proj}_{\mathcal{C}^\circ} w\|_2 > 1$. We define

$$w_0 := \frac{\text{proj}_{\mathcal{C}^\circ} w}{\|\text{proj}_{\mathcal{C}^\circ} w\|_2}.$$

Using characterization (7), we aim at proving that $w_0 = \text{proj}_{\mathcal{C}^\circ \cap \mathcal{B}} w$ (see Figure 4), which would prove (8). First, w_0 belongs to $\mathcal{C}^\circ \cap \mathcal{B}$ by definition. Let $w' \in \mathcal{C}^\circ \cap \mathcal{B}$. For short, denote $w_1 = \text{proj}_{\mathcal{C}^\circ} w$.

$$\begin{aligned} \langle w - w_0, w' - w_0 \rangle &= \langle w - w_1 + w_1 - w_0, w' - w_0 \rangle \\ &= \langle w - w_1, w' - w_0 \rangle + \langle w_1 - w_0, w' - w_0 \rangle \\ &= \frac{1}{\|w_1\|} \langle w - w_1, \|w_1\| w' - w_1 \rangle + \langle w_1 - w_0, w' - w_0 \rangle. \end{aligned}$$




Figure 4: In the case where $\|\text{proj}_{C^\circ} w\|_2 > 1$, we have $w_0 = \text{proj}_{C^\circ \cap B} w$.

The first dot product above is nonpositive by characterization of $w_1 = \text{proj}_{C^\circ} w$, because $\langle w_1, w' \rangle \leq 0$ for $w' \in C^\circ$. Let us prove that the second dot product is also nonpositive. For all $w'' \in C^\circ \cap B$,

$$\|w_1 - w''\| \geq \left| \|w_1\| - \|w''\| \right| \geq \|w_1\| - 1 = \|w_1 - w_0\|,$$

which means that $w_0 = \text{proj}_{C^\circ \cap B} w_1$. Thus, $\langle w_1 - w_0, w' - w_0 \rangle \leq 0$. Therefore,

$$\langle w - w_0, w' - w_0 \rangle \leq 0$$

and (8) is proved. ■

We can now recover via Theorem 10 the classic guarantee for Blackwell's algorithm in the case where the vector payoffs are bounded with respect to the Euclidean norm.

Theorem 27 *Let $M > 0$. Assume that $\|r(a, b)\|_2 \leq M$ for all $a \in \mathcal{A}$ and $b \in \mathcal{B}$. Then, against any sequence of actions $(b_t)_{t \geq 1}$ chosen by the Environment, Blackwell's algorithm guarantees:*

$$\forall T \geq 1, \quad d_2(\bar{r}_T, \mathcal{C}) \leq \frac{2\sqrt{2}M}{\sqrt{T}},$$

where d_2 denotes the Euclidean distance.

Proof With notation from Theorem 26, we have $\max_{x \in \mathcal{X}} h_2(x) - \min_{x \in \mathcal{X}} h_2(x) = 1/2$, and h_2 is 1-strongly convex with respect to $\|\cdot\|_2$ by Proposition 24. According to Theorem 26, Blackwell's algorithm corresponds to the FTRL algorithm associated with h_2 and any sequence of parameters $(\eta_t)_{t \geq 1}$. We can therefore apply Theorem 10 with $\Delta = 1/2$ and $K = 1$, together with Proposition 4 and the result follows. ■

Appendix E. Model with mixed actions

We here present a variant of the model from Section 2.1, in which the decision maker has a finite set of *pure actions* $\mathcal{I} = \{1, \dots, d\}$ from which he is allowed to choose at random. We define the corresponding FTRL algorithms, and state guarantees in expectation, with high probability, and almost-surely. Let the simplex $\Delta_d = \left\{x \in \mathbb{R}_+^d, \sum_{i=1}^d x_i = 1\right\}$ be the set of *mixed actions* (which we identify to the set of probability distributions over \mathcal{I}), \mathcal{B} a set of actions for the Environment, and $r: \mathcal{I} \times \mathcal{B} \rightarrow \mathbb{R}^d$ a payoff function. We linearly extend the payoff function r in its first variable:

$$r(a, b) := \mathbb{E}_{i \sim a} [r(i, b)] = \sum_{i=1}^d a_i r(i, b), \quad a \in \Delta_d, b \in \mathcal{B}.$$

The game is played as follows. At time $t \geq 1$,

- the Decision Maker chooses mixed action $a_t \in \Delta_d$;
- the Environment chooses action $b_t \in \mathcal{B}$;
- the Decision Maker draws pure action $i_t \sim a_t$;
- the Decision Maker observes vector payoff $r_t := r(i_t, b_t)$.

Denote $(\mathcal{F}_t)_{t \geq 1}$ the filtration where \mathcal{F}_t is generated by

$$(a_1, b_1, i_1, \dots, a_{t-1}, b_{t-1}, i_{t-1}, a_t, b_t).$$

An algorithm for the Decision Maker is a sequence of maps $\sigma = (\sigma_t)_{t \geq 1}$ where $\sigma_t : (\Delta_d \times \mathcal{I} \times \mathcal{V}^*)^{t-1} \rightarrow \Delta_d$ so that action a_t is given by

$$a_t = \sigma_t(a_1, i_1, r_1, \dots, a_{t-1}, i_{t-1}, r_{t-1}), \quad t \geq 1.$$

Regarding the Environment, we assume that its choice of action b_t does not depend on i_t , so that $\mathbb{E}[r(i_t, b_t) | \mathcal{F}_t] = \mathbb{E}_{i \sim a_t} [r(i, b_t)] = r(a_t, b_t)$. In this model, Blackwell's condition writes as follows.

Definition 28 (Blackwell's condition for games with mixed actions) *A closed convex cone \mathcal{C} of the payoff space \mathcal{V}^* is a B-set for the game with mixed actions $(\mathcal{I}, \mathcal{B}, r)$ if*

$$\forall x \in \mathcal{C}^\circ, \exists a(x) \in \Delta_d, \forall b \in \mathcal{B}, \quad \langle r(a(x), b), x \rangle \leq 0.$$

Such an application $a : \mathcal{C}^\circ \rightarrow \Delta_d$ is called a $(\mathcal{I}, \mathcal{B}, r, \mathcal{C})$ -oracle.

We can now define the FTRL algorithms similarly as in Section 3.2. Let \mathcal{C} be a closed convex cone of the payoff space \mathcal{V}^* which is assumed to be a B-set for the game with mixed actions $(\mathcal{I}, \mathcal{B}, r)$, $a : \mathcal{C}^\circ \rightarrow \Delta_d$ a $(\mathcal{I}, \mathcal{B}, r, \mathcal{C})$ -oracle, \mathcal{X} a generator of \mathcal{C}° , h a regularizer on \mathcal{X} , and $(\eta_t)_{t \geq 1}$ a positive sequence. Then, the corresponding algorithm writes, for $t \geq 1$,

$$\begin{aligned} \text{compute } & x_t = \nabla h^* \left(\eta_{t-1} \sum_{s=1}^{t-1} r_s \right) \\ \text{compute } & a_t = a(x_t) \\ \text{draw } & i_t \sim a_t \\ \text{observe } & r_t = r(i_t, b_t). \end{aligned}$$

Theorem 29 *Let $\Delta, M, K > 0$, $\|\cdot\|$ be a norm on \mathcal{V} , and $\|\cdot\|_*$ its dual norm on \mathcal{V}^* . We assume:*

- (i) $\max_{x \in \mathcal{X}} h(x) - \min_{x \in \mathcal{X}} h(x) \leq \Delta$,
- (ii) h is K -strongly convex with respect to $\|\cdot\|$,
- (iii) $\|r(a, b)\|_* \leq M$ for all $a \in \Delta_d$ and $b \in \mathcal{B}$.

Then the above algorithm guarantees, with the choice $\eta_t = \sqrt{\Delta K / M^2 t}$ (for $t \geq 1$), against any sequence of actions $(b_t)_{t \geq 1}$ chosen by the Environment:

$$\forall T \geq 1, \quad \mathbb{E}[I_{\mathcal{X}}^*(\bar{r}_T)] \leq 2M \sqrt{\frac{\Delta}{KT}}.$$

Let $\delta \in (0, 1)$. For all $T \geq 1$, we have with probability higher than $1 - \delta$,

$$I_{\mathcal{X}}^*(\bar{r}_T) \leq \frac{M}{\sqrt{T}} \left(2\sqrt{\frac{\Delta}{K}} + \|\mathcal{X}\| \sqrt{2 \log(1/\delta)} \right).$$

Almost-surely,

$$\limsup_{T \rightarrow +\infty} I_{\mathcal{X}}^*(\bar{r}_T) \leq 0.$$

Proof Like in the proof of Theorem 10, Lemma 9 gives:

$$I_{\mathcal{X}}^*(\bar{r}_T) \leq \frac{1}{T} \left(\sum_{t=1}^T \langle r_t, x_t \rangle + 2M \sqrt{\frac{\Delta T}{K}} \right). \quad (9)$$

Consider $X_t = \langle r_t, x_t \rangle$. Then, $(X_t)_{t \geq 1}$ is a sequence of super-martingale differences with respect to filtration $(\mathcal{F}_t)_{t \geq 0}$:

$$\mathbb{E}[\langle r_t, x_t \rangle | \mathcal{F}_t] = \mathbb{E}[\langle r(i_t, b_t), x_t \rangle | \mathcal{F}_t] = \langle \mathbb{E}[r(i_t, b_t) | \mathcal{F}_t], x_t \rangle = \langle r(a_t, b_t), x_t \rangle \leq 0,$$

because x is a $(\mathcal{I}, \mathcal{B}, r, \mathcal{C})$ -oracle. Therefore,

$$\mathbb{E} \left[\sum_{t=1}^T \langle r_t, x_t \rangle \right] = \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}[\langle r_t, x_t \rangle | \mathcal{F}_t] \right] \leq 0.$$

Injecting this in Equation (9) gives the guarantee in expectation:

$$\mathbb{E}[I_{\mathcal{X}}^*(\bar{r}_T)] \leq 2M \sqrt{\frac{\Delta}{KT}}.$$

We now turn to the high probability bound. Let $\delta \in (0, 1)$. From Equation (9), we deduce that

$$I_{\mathcal{X}}^*(\bar{r}_T) \leq 2M \sqrt{\frac{\Delta}{KT}} + \frac{1}{T} \sum_{t=1}^T X_t.$$

Since we have $|X_t| = |\langle r(i_t, b_t), x_t \rangle| \leq \|r(i_t, b_t)\|_* \|x_t\| \leq M \|\mathcal{X}\|$ for all $t \geq 1$, the Azuma–Hoeffding inequality assures that with probability higher than $1 - \delta$,

$$\frac{1}{T} \sum_{t=1}^T X_t \leq M \|\mathcal{X}\| \sqrt{\frac{2 \log(1/\delta)}{T}}$$

and thus

$$I_{\mathcal{X}}^*(\bar{r}_T) \leq \frac{M}{\sqrt{T}} \left(2\sqrt{\frac{\Delta}{K}} + \|\mathcal{X}\| \sqrt{2 \log(1/\delta)} \right).$$

The almost-sure guarantee follows from a standard Borel–Cantelli argument. ■

Appendix F. An algorithm for arbitrary norm global cost

Let $q' \in (1, 2]$. We consider on $\mathcal{X} = \mathcal{B} \cap \mathcal{C}^\circ$ the $\ell_{q'}$ regularizer introduced in Section 3.1:

$$h_{q'}(x) = \begin{cases} \frac{1}{2} \|x\|_{q'}^2 & \text{if } x \in \mathcal{X} \\ +\infty & \text{otherwise,} \end{cases}, \quad x \in \mathcal{V}.$$

Let $(\eta_t)_{t \geq 1}$ a positive sequence, and a the oracle from Remark 12. The algorithm then writes, for $t \geq 1$,

$$\begin{aligned} \text{compute } x_t &= \nabla h_{q'}^* \left(\eta_{t-1} \sum_{s=1}^{t-1} r_s \right) \\ \text{compute } a_t &= \arg \min_{a \in \Delta_d} \sum_{i=1}^d \max(0, z_{ti} a_i + z'_{ti}), \quad \text{where } (z_t, z'_t) = x_t, \\ \text{observe } r_t &:= r(a_t, b_t). \end{aligned}$$

Theorem 30 (Regret bound for an arbitrary norm cost function) *Let $q' \in (1, 2]$ and $\Delta > 0$ such that $\max_{x \in \mathcal{X}} \frac{1}{2} \|x\|_{q'}^2 \leq \Delta$. Then, the above algorithm with coefficients*

$$\eta_t = d^{1/q'-1} \sqrt{\frac{\Delta(q'-1)}{t}}, \quad t \geq 1,$$

guarantees, against any sequence $(\ell_t)_{t \geq 1}$ in $[0, 1]^d$ chosen by the Environment,

$$\forall T \geq 1, \quad \overline{\text{Reg}}_T \leq 2 d^{1-1/q'} \sqrt{\frac{\Delta}{(q'-1)T}}.$$

Proof We aim at applying Theorem 10. According to Proposition 24, because $q' \in (1, 2]$, regularizer $h_{q'}$ is $(q'-1)/d^{2(1-1/q')}$ strongly-convex with respect to $\|\cdot\|_1$. Besides, the payoff function r is bounded by 1 with respect to $\|\cdot\|_\infty$. Indeed, for all $a \in \Delta_d$ and $\ell \in [0, 1]^d$,

$$\|r(a, \ell)\|_\infty = \|(a \odot \ell, \ell)\|_\infty = \max(\|a \odot \ell\|_\infty, \|\ell\|_\infty) \leq 1.$$

And because $0 \in \mathcal{X}$, we have the difference between the highest and the lowest values of $h_{q'}$ on its domain bounded from above as

$$\max_{x \in \mathcal{X}} h_{q'}(x) - \min_{x \in \mathcal{X}} h_{q'}(x) = \max_{x \in \mathcal{X}} \frac{1}{2} \|x\|_{q'}^2 - \min_{x \in \mathcal{X}} \frac{1}{2} \|x\|_{q'}^2 = \max_{x \in \mathcal{X}} \|x\|_{q'}^2 \leq \Delta.$$

Therefore, applying Theorem 10 with $K = (q' - 1)/d^{2(1-1/q')}$, $M = 1$ and norm $\|\cdot\|_1$, together with Proposition 13, gives the result. \blacksquare

Appendix G. Online combinatorial optimization

We illustrate the flexibility of our general framework by giving an alternative construction of an optimal algorithm in the the online combinatorial optimization problem with full information feedback. It is a regret minimization problem in which the actions and the payoffs have a particular structure. Numerous papers were written on the topic, including Gentile and Warmuth (1998); Kivinen and Warmuth (2001); Grove et al. (2001); Takimoto and Warmuth (2003); Kalai and Vempala (2005); Warmuth and Kuzmin (2008); Helmbold and Warmuth (2009); Hazan et al. (2010). A minimax optimal algorithm was given in Koolen et al. (2010). We give below an alternative construction of such an algorithm.

Let $d, m \geq 1$ be integers. Let $\mathcal{I} = \{1, \dots, d\}$ be a finite set. The set of pure actions of the Decision Maker is a set P which contains subsets of \mathcal{I} of cardinality m . Denote $\Delta(P)$ the unit simplex in \mathbb{R}^P and let it be the set of mixed actions by identifying it to the set of probability distributions over P . The game is played as follows. At time $t \geq 1$, the Decision Maker

- chooses mixed action $a_t \in \Delta(P)$;
- draws pure action $p_t \sim a_t$;
- observes payoff vector $v_t \in \mathbb{R}^d$;
- gets payoff $\sum_{i \in p_t} v_{ti}$.

We assume that the choice by the Environment of payoff vector $v_t \in \mathbb{R}^d$ does not depend on pure action p_t . The quantity to minimize is the following regret:

$$\text{Reg}_T = \max_{p \in P} \sum_{t=1}^T \sum_{i \in p} v_{ti} - \sum_{t=1}^T \sum_{i \in p_t} v_{ti}.$$

This problem can be seen as a basic regret minimization problem with pure action set P , and payoff vectors $(\sum_{i \in p} v_i)_{p \in P}$ which belong to $[-m, m]^P$ as soon as we assume $v \in [-1, 1]^d$. The classical Exponential Weights Algorithm (Cesa-Bianchi, 1997) would then guarantee a regret bound of order $m\sqrt{T \log |P|}$. However, our goal is to take advantage of the structure of the problem and to construct a algorithm which guarantees a significantly tighter regret bound, of order $m\sqrt{T \log(d/m)}$, which is known to be minimax optimal (Koolen et al., 2010). To do so, we reduce this problem to a well-chosen approachability game (with mixed actions, as in Section E), which we now present.

Let A be the $d \times |P|$ matrix defined by $A = (\mathbb{1}_{\{i \in p\}})_{\substack{i \in \mathcal{I} \\ p \in P}}$, and for each $p \in P$, let $e_p = (\mathbb{1}_{\{i \in p\}})_{i \in \mathcal{I}} \in \mathbb{R}^d$. Let P (resp. $\Delta(P)$) be the set of pure (resp. mixed) actions for the Decision Maker, $\mathcal{B} = [-1, 1]^d$ the set of actions for the Environment, and consider the following payoff function:

$$r(p, v) = v - \frac{\langle v, e_p \rangle}{m} \mathbf{1} \in \mathbb{R}^d, \quad p \in P, \quad v \in [-1, 1]^d,$$

where $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^d$. The payoff space is therefore $\mathcal{V}^* = \mathbb{R}^d$. The linear extension of the payoff function in its first variable writes

$$r(a, v) = v - \frac{\langle v, Aa \rangle}{m} \mathbf{1}, \quad a \in \Delta(P), \quad v \in [-1, 1]^d.$$

We now choose the generator: let $\mathcal{X} = A(\Delta(P))$ be the image of the simplex $\Delta(P)$ via A seen as a linear map from \mathbb{R}^P to \mathbb{R}^d . Its properties are gathered in the following proposition. In particular, property (v) demonstrates that this choice of \mathcal{X} makes $I_{\mathcal{X}}^*(\bar{r}_T)$ equal to the above regret.

Proposition 31 (i) \mathcal{X} is the convex hull of the points e_p ($p \in P$).

(ii) $\mathcal{X} \subset m\Delta_d$.

(iii) $\|\mathcal{X}\|_1 = m$.

(iv) \mathcal{X} is a generator of $\mathcal{X}^{\circ\circ} = A(\Delta(P))^{\circ\circ}$.

(v) Let $(p_t)_{t \geq 1}$ be a sequence of pure actions chosen by the Decision Maker and $(v_t)_{t \geq 1}$ a sequence of actions chosen by the Environment, and denote $r_t = r(p_t, v_t)$ for all $t \geq 1$ the corresponding payoffs. Then, for all $T \geq 1$,

$$I_{\mathcal{X}}^*(\bar{r}_T) = \frac{1}{T} \text{Reg}_T = \frac{1}{T} \left(\max_{p \in P} \sum_{t=1}^T \sum_{i \in p} v_{ti} - \sum_{t=1}^T \sum_{i \in p_t} v_{ti} \right).$$

Proof By definition, \mathcal{X} is the image of simplex $\Delta(P)$ via linear map A . It is therefore the convex hull of the image by A of the extreme points of $\Delta(P)$. And for $p_0 \in P$, $A(\mathbb{1}_{\{p=p_0\}})_{p \in P} = e_{p_0}$. Hence (i). Each point e_p clearly belongs to $m\Delta_d$, and (ii) is true by convexity of $m\Delta_d$. For each element $x \in m\Delta_d$, we have $\|x\|_1 = m$, which implies (iii). \mathcal{X} is a nonempty convex compact set thanks to (i); Proposition 2 gives (iv). As for the relation

(v), we denote A^* the transpose of A and write

$$\begin{aligned}
 \max_{p \in P} \sum_{t=1}^T \sum_{i \in p} v_{ti} - \sum_{t=1}^T \sum_{i \in p_t} v_{ti} &= \max_{p \in P} \sum_{t=1}^T ((A^* v_t)_p - (A^* v_t)_{p_t}) \\
 &= \max_{a \in \Delta(P)} \sum_{t=1}^T \left(\langle A^* v_t, a \rangle - \left\langle A^* v_t, (\mathbb{1}_{\{p=p_t\}})_{p \in P} \right\rangle \right) \\
 &= \max_{a \in \Delta(P)} \sum_{t=1}^T \left(\langle v_t, Aa \rangle - \left\langle v_t, A (\mathbb{1}_{\{p=p_t\}})_{p \in P} \right\rangle \right) \\
 &= \max_{x \in A(\Delta(P))} \sum_{t=1}^T (\langle v_t, x \rangle - \langle v_t, e_{p_t} \rangle) \\
 &= \max_{x \in \mathcal{X}} \sum_{t=1}^T \left\langle v_t - \frac{\langle v_t, e_{p_t} \rangle}{m} \mathbf{1}, x \right\rangle \\
 &= \max_{x \in \mathcal{X}} \sum_{t=1}^T \langle r(p_t, v_t), x \rangle \\
 &= T \cdot I_{\mathcal{X}}^*(\bar{r}_T),
 \end{aligned}$$

where in the fifth line, we used the fact that for all $x \in \mathcal{X}$, $\langle \mathbf{1}, x \rangle = m$, which is a consequence of (ii). \blacksquare

Proposition 32 $A(\Delta(P))^\circ$ is a B-set for the game with mixed actions $(P, [-1, 1]^d, r)$.

Proof Since \mathcal{X} is a generator of $A(\Delta(P))^\circ$, one can check that the condition that defines a B-set only needs to be verified for $x \in \mathcal{X}$. Let $x \in \mathcal{X}$. By definition of \mathcal{X} , there exists $a \in \Delta(P)$ such that $x = Aa$. Then for $v \in [-1, 1]^d$,

$$\langle r(a, v), x \rangle = \left\langle v - \frac{\langle v, Aa \rangle}{m} \mathbf{1}, Aa \right\rangle = \langle v, Aa \rangle - \langle v, Aa \rangle = 0,$$

which proves the result. \blacksquare

As a consequence of Proposition 31, a point $x \in \mathcal{X}$ only has nonnegative components. We can therefore define

$$h(x) = \begin{cases} \sum_{i=1}^d \frac{x_i}{m} \log \frac{x_i}{m} & \text{for } x \in \mathcal{X} \\ +\infty & \text{otherwise.} \end{cases}$$

Proposition 33 (i) h is a regularizer on \mathcal{X} ;

(ii) $\max_{x \in \mathcal{X}} h - \min_{x \in \mathcal{X}} h(x) \leq \log(d/m)$;

(iii) h is $1/m^2$ -strongly convex with respect to $\|\cdot\|_1$.

Proof For $x \in \mathcal{X} \subset m\Delta_d$, we can write $h(x) = h_{\text{ent}}(x/m) < +\infty$. The 1-strong convexity of h_{ent} with respect to $\|\cdot\|_1$ implies the $1/m^2$ -strong convexity of h with respect to $\|\cdot\|_1$ and (iii) is proved. In particular, h is strictly convex. Besides, the domain of h is \mathcal{X} by definition and (i) is proved. As for (ii), h being convex, its maximum is attained at one of the extreme points e_p ($p \in P$) of \mathcal{X} :

$$\max_{x \in \mathcal{X}} h(x) = \max_{p \in P} h(e_p) = \max_{p \in P} \sum_{i \in p} \frac{1}{m} \log \frac{1}{m} = -\log m.$$

As for the minimum,

$$\min_{x \in \mathcal{X}} h(x) \geq \min_{x \in m\Delta_d} \sum_{i=1}^d \frac{x_i}{m} \log \frac{x_i}{m} = \min_{x \in \Delta_d} \sum_{i=1}^d x_i \log x_i = -\log d.$$

Therefore, $\max_{x \in \mathcal{X}} h - \min_{x \in \mathcal{X}} h(x) \leq -\log m + \log d = \log(d/m)$. ■

We can now consider the FTRL algorithm associated with regularizer h , a $(P, [-1, 1]^d, r, A(\Delta(P))^\circ)$ -oracle a , and a positive sequence of parameters $(\eta_t)_{t \geq 1}$, for $t \geq 1$,

$$\begin{aligned} \text{compute } & x_t = \arg \max_{x \in \mathcal{X}} \left\{ \left\langle \eta_{t-1} \sum_{s=1}^{t-1} r_s, x \right\rangle - h(x) \right\} \\ \text{choose } & a_t = a(x_t) \\ \text{draw } & p_t \sim a_t \\ \text{observe } & r_t = r(p_t, v_t) = v_t - \frac{\langle v_t, A e_{p_t} \rangle}{m} \mathbf{1}. \end{aligned}$$

Theorem 34 *Against any sequence $(v_t)_{t \geq 1}$ in $[-1, 1]^d$ chosen by the Environment, the above algorithm with parameters $\eta_t = \sqrt{\log(d/m)/4m^2 t}$ (for $t \geq 1$) guarantees*

$$\mathbb{E}[\text{Reg}_T] \leq 4m\sqrt{T \log(d/m)}.$$

For $\delta \in (0, 1)$, we have with probability higher than $1 - \delta$,

$$\text{Reg}_T \leq 2m\sqrt{T} \left(2\sqrt{\log(d/m)} + \sqrt{2 \log(1/\delta)} \right).$$

Almost-surely,

$$\limsup_{T \rightarrow +\infty} \frac{1}{T} \text{Reg}_T \leq 0.$$

Proof For all $v \in [-1, 1]^d$ and $p \in P$,

$$\|r(p, v)\|_\infty = \left\| v - \frac{\langle v, A e_p \rangle}{m} \mathbf{1} \right\|_\infty \leq \|v\|_\infty + \frac{\|\mathbf{1}\|_\infty}{m} \sum_{i \in p} |v_i| \leq 2.$$

The result then follows from Theorem 29 applied with $M = 2$, $K = 1/m^2$, the properties of the regularizer h given by Proposition 33, and the relation (v) from Proposition 31. ■

Appendix H. Internal and swap regret

We further illustrate the generality of our framework by recovering the best known algorithms for internet and swap regret minimization. The notion of *internal regret* was introduced by Foster and Vohra (1997). It is an alternative quantity to the usual regret. Foster and Vohra (1997) first established the existence of algorithms which guarantees that the average internal regret is asymptotically nonpositive (see also Fudenberg and Levine (1995, 1999); Hart and Mas-Colell (2000, 2001); Stoltz and Lugosi (2005)). Blum and Mansour (2005) introduced the swap regret, which generalizes both the internal and the basic regret. The optimal bound on the swap regret is known since Blum and Mansour (2005); Stoltz and Lugosi (2005). Later, Perchet (2015) proposed an approachability-based optimal algorithm. We present below the construction of an algorithm similar to Stoltz and Lugosi (2005); Perchet (2015) using the tools introduced in Sections 2 and 3. The internal regret is mentioned at the end of the section as a special case.

The set of pure actions of the Decision Maker is $\mathcal{I} = \{1, \dots, d\}$. At time $t \geq 1$, the Decision Maker

- chooses mixed action $a_t \in \Delta_d$;
- draws pure action $i_t \sim a_t$;
- observes payoff vector $v_t \in \mathbb{R}^d$.

Let Φ be a nonempty subset of $\mathcal{I}^{\mathcal{I}}$. The quantity to minimize is the Φ -regret defined by

$$\text{Reg}_T^\Phi = \max_{\varphi \in \Phi} \sum_{t=1}^T v_{t\varphi(i_t)} - \sum_{t=1}^T v_{ti_t},$$

and can be interpreted as follows. For a given map $\varphi \in \Phi$, $\sum_{t=1}^T v_{t\varphi(i_t)}$ is the cumulative payoff that the Decision Maker would have obtained if he had played pure action $\varphi(i)$ each time he has actually played i (for all $i \in \mathcal{I}$). The Φ -regret therefore compares the actual cumulative payoff of the Decision Maker with the best such quantity (for $\varphi \in \Phi$) in hindsight. The goal is to construct an algorithm which guarantees on the Φ -regret a bound of order $\sqrt{T \log |\Phi|}$. To do so, we reduce this problem to a well-chosen approachability game (with mixed actions as in Section E), which we now present.

Let \mathcal{I} (resp. Δ_d) be the set of pure (resp. mixed) actions for the Decision Maker and $[-1, 1]^d$ the set of actions for the Environment. Let the payoff space be $\mathcal{V}^* = \mathbb{R}^\Phi$ and the target set be \mathbb{R}_-^Φ . We choose the following payoff function:

$$r(i, v) = (v_{\varphi(i)} - v_i)_{\varphi \in \Phi} \in \mathbb{R}^\Phi, \quad i \in \mathcal{I}, v \in [-1, 1]^d.$$

The linear extension of the payoff function in its first variable is

$$r(a, v) = \left(\sum_{i \in \mathcal{I}} a_i (v_{\varphi(i)} - v_i) \right)_{\varphi \in \Phi}, \quad a \in \Delta_d, v \in \mathbb{R}^d.$$

Proposition 35 \mathbb{R}_-^Φ is a B-set for the game with mixed actions $(\mathcal{I}, [-1, 1]^d, r)$.

Proof Let $x = (x_\varphi)_{\varphi \in \Phi} \in (\mathbb{R}_-^\Phi)^\circ = \mathbb{R}_+^\Phi$. Let us prove that there exists $a \in \Delta(\mathcal{I})$ such that for all $v \in [-1, 1]^d$, $\langle r(a, v), x \rangle \leq 0$. First, the property is trivially true if $x = 0$. We assume from now on that $x \neq 0$.

Denote

$$\tilde{x}_{ij} = \sum_{\substack{\varphi \in \Phi \\ \varphi(i)=j}} x_\varphi, \quad i, j \in \mathcal{I}$$

and let us first prove that there exists $a \in \Delta(\mathcal{I})$ such that:

$$\sum_{i \in \mathcal{I}} a_i \tilde{x}_{ij} = a_j \sum_{i \in \mathcal{I}} \tilde{x}_{ji}, \quad j \in \mathcal{I}. \quad (10)$$

Notice that for all $i \in \mathcal{I}$ we have

$$\sum_{j \in \mathcal{I}} \tilde{x}_{ij} = \sum_{j \in \mathcal{I}} \sum_{\substack{\varphi \in \Phi \\ \varphi(i)=j}} x_\varphi = \sum_{\varphi \in \Phi} x_\varphi = \|x\|_1.$$

x being nonzero, the above quantity is also nonzero and the $d \times d$ matrix $(\tilde{x}_{ij}/\|x\|_1)_{i,j \in \mathcal{I}}$ is stochastic and therefore has an invariant measure $a \in \Delta(\mathcal{I})$:

$$\sum_{i \in \mathcal{I}} a_i \frac{\tilde{x}_{ij}}{\|x\|_1} = a_j, \quad j \in \mathcal{I}.$$

Multiplying on both sides by $\|x\|_1$, we get Equation(10):

$$\sum_{i \in \mathcal{I}} a_i \tilde{x}_{ij} = a_j \|x\|_1 = a_j \sum_{i \in \mathcal{I}} \sum_{\substack{\varphi \in \Phi \\ \varphi(j)=i}} x_\varphi = a_j \sum_{i \in \mathcal{I}} \tilde{x}_{ji}, \quad j \in \mathcal{I}.$$

Let $v \in [-1, 1]^d$ and compute $\langle r(a, v), x \rangle$:

$$\begin{aligned} \langle r(a, v), x \rangle &= \sum_{\varphi \in \Phi} x_\varphi \left(\sum_{i \in \mathcal{I}} a_i (v_{\varphi(i)} - v_i) \right) = \sum_{i,j \in \mathcal{I}} a_i (v_j - v_i) \sum_{\substack{\varphi \in \Phi \\ \varphi(i)=j}} x_\varphi \\ &= \sum_{i,j \in \mathcal{I}} a_i (v_j - v_i) \tilde{x}_{ij} = \sum_{j \in \mathcal{I}} v_j \sum_{i \in \mathcal{I}} a_i \tilde{x}_{ij} - \sum_{i,j \in \mathcal{I}} a_i v_i \tilde{x}_{ij} \\ &= \sum_{j \in \mathcal{I}} v_j a_j \sum_{i \in \mathcal{I}} \tilde{x}_{ji} - \sum_{i,j \in \mathcal{I}} a_i v_i \tilde{x}_{ij} = 0, \end{aligned}$$

where we used Equation (10) for the fifth equality. In particular, $\langle r(a, v), x \rangle \leq 0$ and \mathbb{R}_-^Φ is indeed a B-set for the game with mixed actions $(\mathcal{I}, [-1, 1]^d, r)$. \blacksquare

As for the generator, we choose $\mathcal{X} = \Delta(\Phi)$ which is a generator of $(\mathbb{R}_-^\Phi)^\circ$ thanks to Proposition 2. Then the support function of $\Delta(\Phi)$ evaluated at the average payoff is equal to the average Φ -regret:

$$\begin{aligned} I_{\Delta(\Phi)}^*(\bar{r}_T) &= \frac{1}{T} I_{\Delta(\Phi)}^* \left(\sum_{t=1}^T r(i_t, v_t) \right) = \frac{1}{T} \max_{x \in \Delta(\Phi)} \left\langle \sum_{t=1}^T (v_{t\varphi(i_t)} - v_{ti_t})_{\varphi \in \Phi}, x \right\rangle \\ &= \frac{1}{T} \max_{\varphi \in \Phi} \sum_{t=1}^T (v_{t\varphi(i_t)} - v_{ti_t}) = \frac{1}{T} \left(\max_{\varphi \in \Phi} \sum_{t=1}^T v_{t\varphi(i_t)} - \sum_{t=1}^T v_{ti_t} \right) = \frac{1}{T} \text{Reg}_T^\Phi. \end{aligned}$$

On the simplex $\Delta(\Phi)$, we choose the entropic regularizer presented in Section 3.1:

$$h_{\text{ent}}(x) = \begin{cases} \sum_{\varphi \in \Phi} x_\varphi \log x_\varphi & \text{if } x \in \Delta(\Phi) \\ +\infty & \text{otherwise.} \end{cases}$$

Then, the algorithm associated with regularizer h_{ent} , a $(\mathcal{I}, [-1, 1]^d, r, \mathbb{R}_-^\Phi)$ -oracle a and a sequence of positive parameters $(\eta_t)_{t \geq 1}$ is the following. For $t \geq 1$,

$$\begin{aligned} \text{compute } x_{t\varphi} &= \frac{\exp\left(\eta_{t-1} \sum_{s=1}^{t-1} r_{s\varphi}\right)}{\sum_{\varphi' \in \Phi} \exp\left(\eta_{t-1} \sum_{s=1}^{t-1} r_{s\varphi'}\right)}, \quad \varphi \in \Phi \\ \text{choose } a_t &= a(x_t) \\ \text{draw } i_t &\sim a_t \\ \text{observe } r_t &= r(i_t, v_t) = (v_{t\varphi(i_t)} - v_{ti_t})_{\varphi \in \Phi}. \end{aligned}$$

The expression of x_t is explicit and straightforward and the computation of mixed action $a_t = a(x_t)$ via oracle a consists, as shown in the proof of Proposition 35, in finding an invariant measure of a $d \times d$ stochastic matrix, which can be done efficiently. However, the computation of x_t requires to work with $|\Phi|$ components, which can be up to d^d . The algorithm from Blum and Mansour (2005) is much more efficient computationally as its computational cost is polynomial in d .

Theorem 36 *Against any sequence $(v_t)_{t \geq 1}$ in $[-1, 1]^d$ chosen by the Environment, the above algorithm with parameters $\eta_t = \sqrt{\log |\Phi| / 4t}$ (for $t \geq 1$) guarantees*

$$\mathbb{E} [\text{Reg}_T^\Phi] \leq 4\sqrt{T \log |\Phi|}.$$

Let $\delta \in (0, 1)$. With probability higher than $1 - \delta$, we have

$$\frac{1}{T} \text{Reg}_T^\Phi \leq \frac{1}{\sqrt{T}} \left(4\sqrt{\log |\Phi|} + 2\sqrt{2 \log(1/\delta)} \right).$$

Almost-surely,

$$\limsup_{T \rightarrow +\infty} \frac{1}{T} \text{Reg}_T^\Phi \leq 0.$$

Proof For every payoff vector $v \in [-1, 1]^d$ and pure action $i \in \mathcal{I}$,

$$\|r(i, v)\|_\infty = \|(v_{\varphi(i)} - v_i)_{\varphi \in \Phi}\|_\infty \leq 2.$$

The result then follows from Theorem 29 applied with $M = 2$, $K = 1$ and the properties of regularizer h_{ent} given by Proposition 22. \blacksquare

An important special case is when Φ is the set of all transpositions of \mathcal{I} , in other words, the set of maps $\varphi : \mathcal{I} \rightarrow \mathcal{I}$ such that there exists $i \neq j$ in \mathcal{I} such that

$$\varphi(i) = j, \quad \varphi(j) = i, \quad \text{and} \quad \varphi(k) = k \text{ for all } k \notin \{i, j\}.$$

The Φ -regret is then called the *internal regret* and can be written

$$\max_{i, j \in \mathcal{I}} \sum_{t=1}^T \mathbb{1}_{\{i_t = i\}} (v_{tj} - v_{ti}).$$

Since $|\Phi| = d(d-1)$ in this case, Theorem 36 assures that the corresponding algorithm guarantees a bound on the internal regret of order $\sqrt{T \log d}$.