

Set-Valued Approachability and Online Learning with Partial Monitoring

Shie Mannor

*Israel Institute of Technology (Technion)
Faculty of Electrical Engineering
32000 Haifa, Israel*

SHIE@EE.TECHNION.AC.IL

Vianney Perchet

*Université Paris Diderot, LPMA
8 place FM/13
75013 Paris, France*

VIANNEY.PERCHET@NORMALESUP.ORG

Gilles Stoltz

*GREGHEC: HEC Paris — CNRS
1 rue de la Libération
78351 Jouy-en-Josas, France*

STOLTZ@HEC.FR

Editor: Alexander Rakhlin

Abstract

Approachability has become a standard tool in analyzing learning algorithms in the adversarial online learning setup. We develop a variant of approachability for games where there is ambiguity in the obtained reward: it belongs to a set rather than being a single vector. Using this variant we tackle the problem of approachability in games with partial monitoring and develop a simple and generally efficient strategy (i.e., with constant per-step complexity) for this setup. As an important example, we instantiate our general strategy to the case when external regret or internal regret is to be minimized under partial monitoring.

Keywords: online learning, approachability, regret, partial monitoring

1. Introduction

Blackwell's approachability theory and its variants have become a standard and useful tool in analyzing online learning algorithms (Cesa-Bianchi and Lugosi, 2006) and algorithms for learning in games (Hart and Mas-Colell, 2000, 2001). The first application of Blackwell's approachability to learning in the online setup is due to Blackwell (1956b) himself. Numerous other contributions are summarized in the monograph by Cesa-Bianchi and Lugosi (2006). Blackwell's approachability theory enjoys a natural geometric interpretation that allows it to be used in situations where other learning methods (online convex optimization or exponential weights) do not seem to be easily applicable. In some sense, it can be used to go beyond the minimization of the regret to control quantities of a different flavor. Examples of such uses can be found in Mannor et al. (2009), which minimizes the regret together with path constraints, and in Mannor and Shimkin (2008), which minimizes the regret in games whose stage duration is not fixed. Recently, it has been shown by Abernethy et al.

(2011) that approachability and low regret learning are equivalent in the sense that efficient reductions exist from one problem to the other. Another recent paper by Rakhlin et al. (2011) showed that approachability can be analyzed from the perspective of learnability using tools from learning theory.

In this paper we consider approachability and online learning with partial monitoring in games against an arbitrary opponent. That is, we will obtain worst-case performance guarantees: guarantees that are valid for *all* strategies of the opponent. In partial monitoring the decision maker does not know how much reward was obtained and only gets a (random) signal whose distribution depends on the pair of actions taken by the decision maker and the opponent. There are two extremes of this setup that are well studied. On the one extreme we have the case where the signal includes the reward itself (or a signal that can be used to unbiasedly estimate the reward), which is essentially the celebrated bandits setup. The other extreme is the case where the signal is not informative (i.e., it tells the decision maker nothing about the actual reward obtained); this setting then essentially consists of repeating the same situation over and over again, as no information is gained over time. We consider a setup encompassing these situations and more general ones, in which the signal is indicative of the actual reward, but is not necessarily a sufficient statistic thereof. The difficulty is that the decision maker cannot compute the actual reward obtained nor the actions of the opponent.

Regret minimization with partial monitoring (defined in the general sense of Rustichini, 1999) has been studied in several papers in the learning theory community. Piccolboni and Schindelhauer (2001), Mannor and Shimkin (2003), Cesa-Bianchi et al. (2006), Bartók et al. (2010, 2011), Foster and Rakhlin (2012) study games in which an accurate estimation of the rewards (or worst-case rewards) of the decision maker is possible thanks to some statistically sufficient monitoring; in this case, the notion of regret with partial monitoring reduces to the classical notion of regret with full monitoring. A general policy with vanishing external regret with partial monitoring is presented by Lugosi et al. (2008). This policy is based on exponential weights and a specific estimation procedure for the (worst-case) obtained rewards.

In contrast, we devise a general (efficient) algorithm for the problem of approachability under partial monitoring. We then apply it to the more restricted problem of regret minimization. More precisely, we first define a new type of approachability setup, for set-valued functions, which enables to re-derive the extension of approachability to the partial monitoring vector-valued setting proposed by Perchet (2011a). More importantly, we provide concrete algorithms for this approachability problem that are more efficient in the sense that, unlike previous works in the domain, their complexity is constant over all steps. Moreover, their rates of convergence are independent of the game at hand, as in the seminal paper by Blackwell (1956b) but for the first time in this general framework. For example, the recent theoretical study of approachability by Perchet and Quincampoix (2011), which is based on somehow related arguments, does neither provide rates of convergence nor concrete algorithms.

1.1 Outline and Comparison to Known Results

The paper is organized in three main parts. The first part consists of Section 2, where we recall basic facts from approachability theory, when a decision maker faces an arbitrary opponent in the standard vector-valued games setting.

The second part deals with our first contribution, a novel setup for approachability termed “set-valued approachability”, where instead of obtaining a vector-valued reward, the decision maker obtains a set, that represents the ambiguity concerning his reward. In Section 3, we provide a simple characterization of approachable convex sets and an algorithm for the set-valued reward setup under the assumption that the set-valued reward functions are linear. In Section 4 we extend the set-valued approachability setup to problems where the set-valued reward functions are not linear, but rather concave in the mixed action of the decision maker and convex in the mixed action of the opponent. This new concept of set-valued approachability is interesting on its own, as it cannot be directly encompassed into classical vector-valued approachability; yet we retrieve several familiar results (characterization of approachable convex sets, rates of convergence that are independent of the dimension, and so on). More importantly, these results are the key tools for our second series of contributions, which we describe now.

The third part studies approachability in repeated games with partial monitoring. Previous general results in this setup suffered from at least one of the following drawbacks. They were either non-constructive (Rustichini, 1999) or were highly inefficient. The latter drawback refers to strategies that relied on some sort of lifting to the space of probability measures on mixed actions (see e.g., Lehrer and Solan, 2007 and Perchet, 2009, 2011a). They then typically required a fine grid of elements in this lifted space, which had to be progressively refined over time. This construction leads to two main issues: on the one hand, the step complexity continuously increases and becomes prohibitive in the number T of past steps. On the other hand, rates of convergence deteriorate and depend on the dimension. Our aim is therefore to devise algorithms that are efficient (as long as the projection onto some convex set can be done efficiently), with a constant step complexity (although it may depend on parameters of the problem at hand), and with rates of convergence independent of the ambient dimension. Our strategies are the first, to our knowledge, satisfying all of these properties in the general approachability framework. They do so because they do not rely on finer and finer grids; as a byproduct, they can also be considered more natural. Section 5 discusses in greater detail all the points mentioned in this paragraph.

More precisely, we state in Section 5.1 the necessary and sufficient condition for approachability in games with partial monitoring and show in Section 5.2 how to apply set-valued approachability framework to the repeated vector-valued games with partial monitoring. In Section 5.3 we then consider a specific type of games where the signaling structure possesses a special property, called bi-piecewise linearity, that can be exploited to derive simple, constructive and efficient strategies. This type of games is rich enough as it encompasses several useful special cases discussed in the later sections. In Section 5.4, we mention the general signaling case and explain how it is possible to approach certain special sets such as polytopes efficiently (thanks to a reduction to bi-piecewise linearity) with the same dimension-independent rates of convergence—and even general convex sets, although inefficiently in the latter case.

As an important other example of a setting where bi-piecewise linearity holds, we apply in Section 6 the results of Section 5.3 to both external-regret and internal-regret minimization in repeated games with partial monitoring. In this specific case, our algorithms have rates similar to the ones obtained by Lugosi et al. (2008) but slower than Perchet (2011b); however our proof is direct and simpler and the strategy is efficient.

1.2 Mixed Actions versus Pure Actions

Most of Sections 2–4 (classical approachability and set-valued approachability) is concerned with mixed actions, while Sections 5–6 (approachability in games with partial monitoring) are focused on pure actions. The explanation for this is as follows. Even though pure actions are inherent to the model of partial monitoring, the reduction from approachability in games with partial monitoring to set-valued approachability, as described in Section 5.2, is to set-valued approachability with mixed actions.

2. Some Basic Facts from Approachability Theory

In this section we recall the most basic version of Blackwell’s approachability theorem for vector-valued payoff functions.

We consider a vector-valued game between two players, a decision maker (first player) and an opponent (second player), with respective finite action sets \mathcal{A} and \mathcal{B} , whose cardinalities are referred to as $N_{\mathcal{A}}$ and $N_{\mathcal{B}}$. We denote by d the dimension of the reward vectors and equip \mathbb{R}^d with the ℓ^2 -norm $\|\cdot\|_2$. The payoff function of the first player is given by a mapping $m : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}^d$, which is multi-linearly extended to $\Delta(\mathcal{A}) \times \Delta(\mathcal{B})$, the set of product-distributions over $\mathcal{A} \times \mathcal{B}$.

We consider a framework in which mixed actions are taken. We denote by $\mathbf{x}_1, \mathbf{x}_2, \dots$ and $\mathbf{y}_1, \mathbf{y}_2, \dots$ the actions in $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$ sequentially taken by each player. We assume a full or bandit monitoring for the first player: at the end of round t , when receiving the payoff $m(\mathbf{x}_t, \mathbf{y}_t)$, either the mixed action \mathbf{y}_t (full monitoring) or only the indicated payoff (bandit monitoring) is revealed to him.

Strategies of the players are defined as mappings associating the information available at the beginning of each round $t \geq 1$ with a mixed action. In particular, strategies of the first player in the case of full monitoring associate with $\mathbf{x}_1, \dots, \mathbf{x}_{t-1}$ and $\mathbf{y}_1, \dots, \mathbf{y}_{t-1}$ a mixed action $\mathbf{x}_t \in \Delta(\mathcal{A})$, while in the case of bandit monitoring, they do this association based on $\mathbf{x}_1, \dots, \mathbf{x}_{t-1}$ and $m(\mathbf{x}_1, \mathbf{y}_1), \dots, m(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})$. We do not restrict the opponent and assume a full monitoring for him: his strategies associate with $\mathbf{x}_1, \dots, \mathbf{x}_{t-1}$ and $\mathbf{y}_1, \dots, \mathbf{y}_{t-1}$ a mixed action $\mathbf{y}_t \in \Delta(\mathcal{B})$.

2.1 Necessary and Sufficient Condition for Approachability

Given a set \mathcal{C} , the aim of the first player is to ensure that his average payoff converges to \mathcal{C} , while the second player wants to prevent it. This gives rise to Blackwell’s classical definition of approachability. (Here, we state it as m -approachability to remind the reader, in the notation, that the underlying payoff function is m .)

Definition 1 Given a function $m : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}^d$, a set $\mathcal{C} \subseteq \mathbb{R}^d$ is m -approachable by the first player if he has a strategy such that, for all $\varepsilon > 0$, there exists an integer T_ε such that for all strategies of the second player,

$$\mathbb{P} \left\{ \forall T \geq T_\varepsilon, \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T m(\mathbf{x}_t, \mathbf{y}_t) \right\|_2 \leq \varepsilon \right\} \geq 1 - \varepsilon.$$

In particular, the first player has a strategy that ensures that the average of his vector-valued payoffs converges almost surely to the set \mathcal{C} , uniformly with respect to the strategies of the second player.

As will be recalled below in Theorem 3, even stronger approachability guarantees can be achieved. Indeed, the first player has deterministic strategies such that, for all (deterministic or randomized) strategies of the second player, with probability 1, for all $T \geq 1$,

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T m(\mathbf{x}_t, \mathbf{y}_t) \right\|_2 \leq \beta(T),$$

where $\beta(\cdot)$ is some decreasing mapping to 0 to be determined later.

For closed convex sets there is a simple characterization of approachability that is a direct consequence of von Neumann's minimax theorem.

Theorem 2 (see Blackwell, 1956a, Theorem 3) A closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ is approachable if and only if

$$\forall \mathbf{y} \in \Delta(\mathcal{B}), \exists \mathbf{x} \in \Delta(\mathcal{A}), \quad m(\mathbf{x}, \mathbf{y}) \in \mathcal{C}.$$

2.2 An Associated Strategy (whose Efficiency Depends on the Geometry of \mathcal{C})

Blackwell suggested a simple strategy with a geometric flavor; it only requires bandit monitoring.

Play an arbitrary \mathbf{x}_1 . For $t \geq 1$, given the vector-valued quantity

$$\widehat{m}_t = \frac{1}{t} \sum_{s=1}^t m(\mathbf{x}_s, \mathbf{y}_s),$$

compute the projection c_t (in ℓ^2 -norm) of \widehat{m}_t on \mathcal{C} . Find a mixed action \mathbf{x}_{t+1} that solves the minimax equation

$$\min_{\mathbf{x} \in \Delta(\mathcal{A})} \max_{\mathbf{y} \in \Delta(\mathcal{B})} \langle \widehat{m}_t - c_t, m(\mathbf{x}, \mathbf{y}) - c_t \rangle, \tag{1}$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product in \mathbb{R}^d .

The stated minimax problem for determining \mathbf{x}_{t+1} can be solved efficiently using, e.g., linear programming: the associated complexity is polynomial in $N_{\mathcal{A}}$ and $N_{\mathcal{B}}$. This strategy is efficient if computing the required projections onto \mathcal{C} in ℓ^2 -norm can be performed efficiently.

The strategy presented above enjoys the following rates of convergence for approachability, which can be derived as a special case of the results stated and proved in Theorem 25 later in this paper.

Theorem 3 (see Blackwell, 1956a, Theorems 1 and 3) *We consider an approachable closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ and we denote by M a bound in norm over m , i.e.,*

$$\max_{(a,b) \in \mathcal{A} \times \mathcal{B}} \|m(a,b)\|_2 \leq M.$$

The above strategy ensures that for all strategies of the second player, with probability 1, for all $T \geq 1$,

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T m(\mathbf{x}_t, \mathbf{y}_t) \right\|_2 \leq \frac{2M}{\sqrt{T}}.$$

3. Set-Valued Approachability for Finite Games

In this section we extend the results from the previous section to set-valued payoff functions in the case of full monitoring. We denote by $\mathcal{S}(\mathbb{R}^d)$ the set of all subsets of \mathbb{R}^d and consider a set-valued payoff function $\bar{m} : \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{S}(\mathbb{R}^d)$. When the players choose respective actions $a \in \mathcal{A}$ and $b \in \mathcal{B}$, the first player gets the subset $\bar{m}(a,b)$ as a payoff. This models the ambiguity or uncertainty associated with some true underlying payoff.

3.1 Mixed Actions Taken and Observed

For the moment, we only consider the case of mixed actions taken and observed, keeping the same definition of a strategy as in the previous section. (The next subsection will briefly explain, for the sake of completeness, how to deal with the case of pure actions taken and observed.)

We extend \bar{m} multi-linearly to $\Delta(\mathcal{A}) \times \Delta(\mathcal{B})$ and even to $\Delta(\mathcal{A} \times \mathcal{B})$, the set of joint probability distributions on $\mathcal{A} \times \mathcal{B}$, as follows. Let

$$\mu = (\mu_{a,b})_{(a,b) \in \mathcal{A} \times \mathcal{B}}$$

be such a joint probability distribution; then $\bar{m}(\mu)$ is defined as a finite convex combination¹ of subsets of \mathbb{R}^d ,

$$\bar{m}(\mu) = \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} \mu_{a,b} \bar{m}(a,b).$$

The product-distribution of two elements $\mathbf{x} = (x_a)_{a \in \mathcal{A}} \in \Delta(\mathcal{A})$ and $\mathbf{y} = (y_b)_{b \in \mathcal{B}} \in \Delta(\mathcal{B})$ will be denoted by $\mathbf{x} \otimes \mathbf{y}$; it gives a probability mass of $x_a y_b$ to each pair $(a,b) \in \mathcal{A} \times \mathcal{B}$. When μ is such a product-distribution, we use the notation $\bar{m}(\mu) = \bar{m}(\mathbf{x}, \mathbf{y})$.

We can now describe how the game proceeds. At each round t , the players choose simultaneously respective mixed actions $\mathbf{x}_t \in \Delta(\mathcal{A})$ and $\mathbf{y}_t \in \Delta(\mathcal{B})$. Full monitoring takes place for the first player: he observes \mathbf{y}_t at the end of round t and he gets the subset $\bar{m}(\mathbf{x}_t, \mathbf{y}_t)$ as a payoff (which, again, accounts for the uncertainty).

1. For two sets S, T and $\alpha \in [0, 1]$, the convex combination $\alpha S + (1 - \alpha)T$ is defined as

$$\{\alpha s + (1 - \alpha)t, \quad s \in S \text{ and } t \in T\}.$$

3.1.1 DEFINITION OF SET-VALUED APPROACHABILITY

We are interested in the behavior of

$$\frac{1}{T} \sum_{t=1}^T \bar{m}(\mathbf{x}_t, \mathbf{y}_t) = \bar{m}(\nu_T), \quad \text{where} \quad \nu_T := \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \otimes \mathbf{y}_t$$

is the empirical joint distribution of mixed actions taken during the first T rounds.

The distance of this set $\bar{m}(\nu_T)$ to the target set \mathcal{C} will be measured in a worst-case sense (à la Hausdorff): we denote by

$$\varepsilon_T = \sup_{\xi \in \bar{m}(\nu_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2$$

the smallest value such that $\bar{m}(\nu_T)$ is included in an ε_T -neighborhood of \mathcal{C} . Approachability of a set \mathcal{C} with the set-valued payoff function \bar{m} then simply means that the sequence of ε_T tends almost-surely to 0, uniformly with respect to the strategies of the second player. This is made formal in the following definition.

Definition 4 *Given a set-valued payoff function $\bar{m} : \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{S}(\mathbb{R}^d)$, a set $\mathcal{C} \subseteq \mathbb{R}^d$ is \bar{m} -approachable by the first player if he has a strategy such that, for all $\varepsilon > 0$, there exists an integer T_ε such that for all strategies of the second player,*

$$\mathbb{P} \left\{ \forall T \geq T_\varepsilon, \sup_{\xi \in \bar{m}(\nu_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2 \leq \varepsilon \right\} \geq 1 - \varepsilon.$$

When the set-valued function \bar{m} is clear from the context, we will simply say that \mathcal{C} is set-valued approachable. Actually, just as in the classical case of approachability, the bounds exhibited below in Theorem 8 will be for deterministic strategies of the first player and will read as follows: for all (deterministic or randomized) strategies of the second player, with probability 1, for all $T \geq 1$,

$$\sup_{\xi \in \bar{m}(\nu_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2 \leq \beta(T),$$

where $\beta(\cdot)$ is a mapping decreasing to 0 to be determined.

3.1.2 A USEFUL CONTINUITY LEMMA

Before proceeding we provide a continuity lemma. It can be reformulated as indicating that for all joint distributions μ and ν over $\mathcal{A} \times \mathcal{B}$, the set $\bar{m}(\mu)$ is contained in an $M \|\mu - \nu\|_1$ -neighborhood of $\bar{m}(\nu)$, where M is a bound in ℓ^2 -norm on \bar{m} . This is a result that we will use repeatedly.

Definition 5 *The set-valued function $\bar{m} : \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{S}(\mathbb{R}^d)$ is bounded in ℓ^2 -norm by M if*

$$\forall (a, b) \in \mathcal{A} \times \mathcal{B}, \quad \sup_{\xi \in \bar{m}(a,b)} \|\xi\|_2 \leq M.$$

Lemma 6 *Let μ and ν be two probability distributions over $\mathcal{A} \times \mathcal{B}$. We assume that the set-valued function \bar{m} is bounded in ℓ^2 -norm by M . Then*

$$\sup_{\xi \in \bar{m}(\mu)} \inf_{c \in \bar{m}(\nu)} \|\xi - c\|_2 \leq M \|\mu - \nu\|_1 \leq M \sqrt{N_{\mathcal{A}} N_{\mathcal{B}}} \|\mu - \nu\|_2,$$

where the norms in the right-hand side are respectively the ℓ^1 and ℓ^2 -norms between probability distributions.

Proof Let ξ be an element of $\bar{m}(\mu)$; it can be written as

$$\xi = \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} \mu_{a,b} \zeta_{a,b}$$

for some elements $\zeta_{a,b} \in \bar{m}(a, b)$. We consider

$$c = \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} \nu_{a,b} \zeta_{a,b},$$

which is an element of $\bar{m}(\nu)$. Then by the triangle inequality,

$$\|\xi - c\|_2 = \left\| \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} (\mu_{a,b} - \nu_{a,b}) \zeta_{a,b} \right\|_2 \leq \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} |\mu_{a,b} - \nu_{a,b}| \|\zeta_{a,b}\|_2 \leq M \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} |\mu_{a,b} - \nu_{a,b}|.$$

This entails the first claimed inequality. The second one follows from an application of the Cauchy-Schwarz inequality. ■

Corollary 7 *If the set-valued function \bar{m} is bounded in norm, then for all $\mathbf{y} \in \Delta(\mathcal{B})$, the mappings $D_{\mathbf{y}} : \Delta(\mathcal{A}) \rightarrow \mathbb{R}$ defined, for all $\mathbf{x} \in \Delta(\mathcal{A})$, by*

$$D_{\mathbf{y}}(\mathbf{x}) = \sup_{\xi \in \bar{m}(\mathbf{x}, \mathbf{y})} \inf_{c \in \mathcal{C}} \|c - \xi\|_2$$

are continuous.

Proof We show that for all $\mathbf{x}, \mathbf{x}' \in \Delta(\mathcal{A})$, the condition $\|\mathbf{x}' - \mathbf{x}\|_1 \leq \varepsilon$ implies that $D_{\mathbf{y}}(\mathbf{x}) - D_{\mathbf{y}}(\mathbf{x}') \leq M\varepsilon$, where M is the bound in ℓ^2 -norm over \bar{m} . Indeed, fix $\delta > 0$ and let $\xi_{\delta, \mathbf{x}} \in \bar{m}(\mathbf{x}, \mathbf{y})$ be such that

$$D_{\mathbf{y}}(\mathbf{x}) \leq \inf_{c \in \mathcal{C}} \|c - \xi_{\delta, \mathbf{x}}\|_2 + \delta. \tag{2}$$

By Lemma 6 (with the choices $\mu = \mathbf{x} \otimes \mathbf{y}$ and $\nu = \mathbf{x}' \otimes \mathbf{y}$),

$$\inf_{\xi' \in \bar{m}(\mathbf{x}', \mathbf{y})} \|\xi_{\delta, \mathbf{x}} - \xi'\| \leq M\varepsilon,$$

and therefore, there exists $\xi_{\delta, \mathbf{x}'} \in \bar{m}(\mathbf{x}', \mathbf{y})$ such that $\|\xi_{\delta, \mathbf{x}} - \xi_{\delta, \mathbf{x}'}\|_2 \leq M\varepsilon + \delta$. The triangle inequality entails that

$$\inf_{c \in \mathcal{C}} \|c - \xi_{\delta, \mathbf{x}}\|_2 \leq \inf_{c \in \mathcal{C}} \|c - \xi_{\delta, \mathbf{x}'}\|_2 + M\varepsilon + \delta.$$

Substituting in (2), we get that

$$D_{\mathbf{y}}(\mathbf{x}) \leq M\varepsilon + 2\delta + \inf_{c \in \mathcal{C}} \|c - \xi_{\delta, \mathbf{x}'}\|_2 \leq M\varepsilon + 2\delta + D_{\mathbf{y}}(\mathbf{x}'),$$

which, letting $\delta \rightarrow 0$, proves our continuity claim. ■

3.1.3 NECESSARY AND SUFFICIENT CONDITION FOR SET-VALUED APPROACHABILITY

This condition will be referred to as (SVAC), an acronym that stands for “set-valued approachability condition.”

Theorem 8 *Suppose that the set-valued function \bar{m} is bounded in norm. A closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ is \bar{m} -approachable if and only if the following set-valued approachability condition is satisfied,*

$$\forall \mathbf{y} \in \Delta(\mathcal{B}), \quad \exists \mathbf{x} \in \Delta(\mathcal{A}), \quad \bar{m}(\mathbf{x}, \mathbf{y}) \subseteq \mathcal{C}. \tag{SVAC}$$

In this case, an \bar{m} -approaching strategy for \mathcal{C} is an m -approaching strategy of $\tilde{\mathcal{C}}$ defined below at (3) and (4). It satisfies, for all $T \geq 1$,

$$\sup_{\xi \in \bar{m}(\nu_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2 \leq 2M \sqrt{\frac{N_{\mathcal{A}} N_{\mathcal{B}}}{T}},$$

where M is a bound in ℓ^2 -norm on \bar{m} .

Proof [of the necessity of Condition (SVAC)] If the condition does not hold, then there exists $\mathbf{y}_0 \in \Delta(\mathcal{B})$ such that for every $\mathbf{x} \in \mathcal{A}$, the set $\bar{m}(\mathbf{x}, \mathbf{y}_0)$ is not included in \mathcal{C} , i.e., it contains at least one point not in \mathcal{C} . We consider the mapping $D_{\mathbf{y}_0}$ defined in the statement of Corollary 7. Since \mathcal{C} is closed, distances of given individual points to \mathcal{C} are achieved; therefore, by the choice of \mathbf{y}_0 , we get that $D_{\mathbf{y}_0}(\mathbf{x}) > 0$ for all $\mathbf{x} \in \Delta(\mathcal{A})$. Now, since $D_{\mathbf{y}_0}$ is continuous on the compact set $\Delta(\mathcal{A})$, as asserted by the indicated corollary, it attains its minimum, whose value we denote by $D_{\min} > 0$.

Assume now that the second player chooses at each round $\mathbf{y}_t = \mathbf{y}_0$ as his mixed action. Then, denoting

$$\bar{\mathbf{x}}_T = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t,$$

we get that $\nu_T = \bar{\mathbf{x}}_T \otimes \mathbf{y}_0$, and hence, for all strategies of the first player and for all $T \geq 1$,

$$\sup_{\xi \in \bar{m}(\nu_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2 = D_{\mathbf{y}_0}(\bar{\mathbf{x}}_T) \geq D_{\min} > 0,$$

which shows that \mathcal{C} is not approachable. ■

We now prove in a constructive way, by exhibiting a suitable strategy (the one alluded at in the statement of the theorem), that (SVAC) is sufficient for set-valued approachability.

We identify probability distributions over $\mathcal{A} \times \mathcal{B}$ with vectors in $\mathbb{R}^{\mathcal{A} \times \mathcal{B}}$ and consider the vector-valued payoff function

$$m : (a, b) \in \mathcal{A} \times \mathcal{B} \mapsto \delta_{(a,b)} \in \mathbb{R}^{\mathcal{A} \times \mathcal{B}}, \tag{3}$$

where $\delta_{(a,b)}$ is the point mass on (a, b) . We extend m to $\Delta(\mathcal{A}) \times \Delta(\mathcal{B})$ in a multi-linear fashion. The target set will be

$$\tilde{\mathcal{C}} = \{ \mu \in \Delta(\mathcal{A} \times \mathcal{B}) : \overline{m}(\mu) \subseteq \mathcal{C} \}. \tag{4}$$

The linearity of the function \overline{m} on $\Delta(\mathcal{A} \times \mathcal{B})$ entails that if \mathcal{C} is a convex set (respectively, a closed set, or a polyhedron), then $\tilde{\mathcal{C}}$ is a convex set as well (respectively, a closed set, or a polyhedron). In the case where $\tilde{\mathcal{C}}$ is a polyhedron, it is actually a polytope (that is, a compact polyhedron).

We then consider the m -approaching strategy of $\tilde{\mathcal{C}}$ described in (1) and now prove that it enjoys the convergence guarantees stated in Theorem 8.

Lemma 9 *Condition (SVAC) is equivalent to the m -approachability of $\tilde{\mathcal{C}}$.*

Proof Since \mathcal{C} and thus $\tilde{\mathcal{C}}$ are closed and convex sets, we can resort to Theorem 2. The latter states that the m -approachability of $\tilde{\mathcal{C}}$ is equivalent to the fact that for all $\mathbf{y} \in \Delta(\mathcal{B})$, there exists some $\mathbf{x} \in \Delta(\mathcal{A})$ such that $\mu = m(\mathbf{x}, \mathbf{y}) = \mathbf{x} \otimes \mathbf{y}$, the product-distribution between \mathbf{x} and \mathbf{y} , belongs to $\tilde{\mathcal{C}}$, i.e., satisfies $\overline{m}(\mu) = \overline{m}(\mathbf{x}, \mathbf{y}) \subseteq \mathcal{C}$. ■

The definition (3) of m entails the rewriting

$$\nu_T = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \otimes \mathbf{y}_t = \frac{1}{T} \sum_{t=1}^T m(\mathbf{x}_t, \mathbf{y}_t).$$

Let $P_{\tilde{\mathcal{C}}}$ denote the projection operator onto $\tilde{\mathcal{C}}$; the quantities at hand in the definition of m -approachability of $\tilde{\mathcal{C}}$ are given by

$$\varepsilon_T = \left\| \nu_T - P_{\tilde{\mathcal{C}}}(\nu_T) \right\|_2 = \inf_{\mu \in \tilde{\mathcal{C}}} \|\nu_T - \mu\|_2.$$

We now relate them to the ones arising in the definition of \overline{m} -approachability of \mathcal{C} .

Lemma 10 *The following upper bound holds,*

$$\sup_{\xi \in \overline{m}(\nu_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2 \leq M \sqrt{N_{\mathcal{A}} N_{\mathcal{B}}} \varepsilon_T.$$

Proof Lemma 6 entails that the sets $\overline{m}(\nu_T)$ are included in $M \sqrt{N_{\mathcal{A}} N_{\mathcal{B}}} \varepsilon_T$ -neighborhoods of $\overline{m}(P_{\tilde{\mathcal{C}}}(\nu_T))$. Since by definition of $\tilde{\mathcal{C}}$, one has $\overline{m}(P_{\tilde{\mathcal{C}}}(\nu_T)) \subseteq \mathcal{C}$, we get in particular that the sets $\overline{m}(\nu_T)$ are included in $M \sqrt{N_{\mathcal{A}} N_{\mathcal{B}}} \varepsilon_T$ -neighborhoods of \mathcal{C} , which is exactly the statement of the lemma. ■

Proof [of the sufficiency of Condition (SVAC)] First, Lemma 9 and Condition (SVAC) entail, via Theorem 3, that the considered strategy m -approaches $\tilde{\mathcal{C}}$, at the following rate: $\varepsilon_T \leq 2/\sqrt{T}$, with probability 1. Second, Lemma 10 indicates that this strategy also \bar{m} -approaches \mathcal{C} , at the stated rate of $2M\sqrt{N_{\mathcal{A}}N_{\mathcal{B}}/T}$, with probability 1. ■

3.1.4 REMARKS: ON EFFICIENCY; ON FULL VERSUS BANDIT MONITORINGS

Note that, as explained around (1), the considered strategy for m -approaching $\tilde{\mathcal{C}}$, or equivalently \bar{m} -approaching \mathcal{C} , is efficient as long as projections in ℓ^2 -norm onto the set $\tilde{\mathcal{C}}$ defined in (4) can be computed efficiently. The latter depends on the respective geometries of \bar{m} and \mathcal{C} . We will provide examples of favorable cases (see, e.g., Section 6.1 about minimization of external regret under partial monitoring). In the sequel the notion of “efficiency up to a projection oracle” will refer to this efficiency depending solely on the efficient computation of the needed projections.

The proposed strategy does not require full monitoring, although it seems to rely on the observation of the pair of played mixed actions $m(\mathbf{x}_t, \mathbf{y}_t)$. With bandit monitoring, only the played sets $\bar{m}(\mathbf{x}_t, \mathbf{y}_t)$ would be available, not the \mathbf{y}_t themselves; in that case, the player can act as if the other player chose any \mathbf{y}'_t that generates this set, i.e., such that $\bar{m}(\mathbf{x}_t, \mathbf{y}'_t) = \bar{m}(\mathbf{x}_t, \mathbf{y}_t)$.

3.2 Pure Actions Taken and Observed

It is well-known that the basic results recalled in Section 2 extend to the case of pure actions. We briefly explain here how the developed theory of set-valued approachability for games with mixed actions extends as well to the case of pure actions taken (still under full monitoring).

The game goes as follows. At each round t , the players choose simultaneously respective pure actions $A_t \in \mathcal{A}$ and $B_t \in \mathcal{B}$, possibly at random according to distributions \mathbf{x}_t and \mathbf{y}_t . As a result, the first player gets the subset $\bar{m}(A_t, B_t)$ as a payoff and observes B_t . Strategies for the players now associate with A_1, \dots, A_{t-1} and B_1, \dots, B_{t-1} mixed actions \mathbf{x}_t and \mathbf{y}_t , according to which A_t and B_t are drawn independently.

We are interested in the behavior of

$$\frac{1}{T} \sum_{t=1}^T \bar{m}(A_t, B_t) = \bar{m}(\pi_T), \quad \text{where} \quad \pi_T := \frac{1}{T} \sum_{t=1}^T \delta_{(A_t, B_t)}$$

is the empirical distribution of the pairs (A_t, B_t) of actions taken during the first T rounds. The definition of set-valued approachability extends as follows.

Definition 11 *A set $\mathcal{C} \subseteq \mathbb{R}^d$ is \bar{m} -approachable by the first player with pure actions if he has a strategy such that, for all $\varepsilon > 0$, there exists an integer T_ε such that for all strategies of the second player,*

$$\mathbb{P} \left\{ \forall T \geq T_\varepsilon, \quad \sup_{\xi \in \bar{m}(\pi_T)} \inf_{c \in \mathcal{C}} \|\xi - c\|_2 \leq \varepsilon \right\} \geq 1 - \varepsilon.$$

The fact that Condition (SVAC) is still a necessary and sufficient condition for \bar{m} -approachability with pure actions of a closed convex set $\mathcal{C} \subseteq \mathbb{R}^d$ (where \bar{m} is bounded) can be seen as follows.

Concerning the proof of the sufficiency of this condition, first recall that Lemma 9 indicates that Condition (SVAC) is equivalent to the m -approachability of $\tilde{\mathcal{C}}$. In view of a version of Theorem 3 for pure actions (e.g., Theorem II.4.3 of Mertens et al., 1994) the strategy described around (3) and (4), with the replacement of the ν_T by the π_T and extra random draws of the pure actions A_t according to the mixed distributions \mathbf{x}_t thus computed, is such that the quantities

$$\varepsilon'_T = \left\| \pi_T - P_{\tilde{\mathcal{C}}}(\pi_T) \right\|_2 = \inf_{\mu \in \tilde{\mathcal{C}}} \|\pi_T - \mu\|_2$$

satisfy the following convergence guarantees. For all $\delta \in (0, 1)$, there exists an integer T_δ such that for all strategies of the second player,

$$\mathbb{P}\{\forall T \geq T_\delta, \varepsilon'_T \leq \delta\} \geq 1 - \delta.$$

This shows that this strategy also \bar{m} -approaches \mathcal{C} since Lemma 10 is valid with the respective replacements of ν_T and ε_T by π_T and ε'_T .

The proof of the necessity of the condition is the same as for mixed actions taken, with the addition of a concentration argument. Indeed, by martingale convergence (e.g., repeated uses of the Hoeffding-Azuma inequality together with an application of the Borel-Cantelli lemma), $\delta_T = \|\pi_T - \nu_T\|_1$ converges to zero almost surely as T goes to infinity. By applying Lemma 6 and by using the notation of the proof of Theorem 8, we get

$$\sup_{\xi \in \bar{m}(\pi_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2 \geq \sup_{\xi \in \bar{m}(\nu_T)} \inf_{c \in \mathcal{C}} \|c - \xi\|_2 - M\delta_T \geq D_{\min} - M\delta_T,$$

and we simply take the \liminf in the above inequalities to conclude the argument.

4. Set-Valued Approachability for Concave–Convex Set-Valued Games

We consider in this section the same setting of mixed actions taken and observed as in Section 3.1, that is, we deal with set-valued payoff functions $\bar{m} : \Delta(\mathcal{A}) \times \Delta(\mathcal{B}) \rightarrow \mathcal{S}(\mathbb{R}^d)$ under full monitoring. However, in the previous section \bar{m} was linear on $\Delta(\mathcal{A}) \times \Delta(\mathcal{B})$, an assumption that we now weaken while still having that (SVAC) is the necessary and sufficient condition for set-valued approachability. The price to pay for this is the loss of the exhibited efficiency (up to a projection oracle) of the approaching strategies and an inferior convergence rate.

Formally, the functions $\bar{m} : \Delta(\mathcal{A}) \times \Delta(\mathcal{B}) \rightarrow \mathcal{S}(\mathbb{R}^d)$ that we will consider will satisfy one or several of the following properties.

Definition 12 *The set-valued function $\bar{m} : \Delta(\mathcal{A}) \times \Delta(\mathcal{B}) \rightarrow \mathcal{S}(\mathbb{R}^d)$ is bounded in ℓ^2 -norm by M if*

$$\forall (\mathbf{x}, \mathbf{y}) \in \Delta(\mathcal{A}) \times \Delta(\mathcal{B}), \quad \sup_{\xi \in \bar{m}(\mathbf{x}, \mathbf{y})} \|\xi\|_2 \leq M.$$

Definition 13 A function $\bar{m} : \Delta(\mathcal{A}) \times \Delta(\mathcal{B}) \rightarrow \mathcal{S}(\mathbb{R}^d)$ is uniformly continuous in its first argument if for every $\varepsilon > 0$, there exists $\eta > 0$ such that for all $\mathbf{x}, \mathbf{x}' \in \Delta(\mathcal{A})$ satisfying $\|\mathbf{x} - \mathbf{x}'\|_1 \leq \eta$ and for all $\mathbf{y} \in \Delta(\mathcal{B})$, the set $\bar{m}(\mathbf{x}', \mathbf{y})$ is included in an ε -neighborhood of $\bar{m}(\mathbf{x}, \mathbf{y})$ in the Euclidean norm. Put differently,

$$\sup_{\xi \in \bar{m}(\mathbf{x}', \mathbf{y})} \inf_{c \in \bar{m}(\mathbf{x}, \mathbf{y})} \|\xi - c\|_2 \leq \varepsilon \quad \text{or} \quad \bar{m}(\mathbf{x}', \mathbf{y}) \subseteq \bar{m}(\mathbf{x}, \mathbf{y}) + \varepsilon \mathbf{B},$$

where \mathbf{B} is the unit Euclidean ball in \mathbb{R}^d .

Uniform continuity in the second argument is defined symmetrically.

Definition 14 A function $\bar{m} : \Delta(\mathcal{A}) \times \Delta(\mathcal{B}) \rightarrow \mathcal{S}(\mathbb{R}^d)$ is concave in its first argument if for all $\mathbf{x}, \mathbf{x}' \in \Delta(\mathcal{A})$, all $\mathbf{y} \in \Delta(\mathcal{B})$, and all $\alpha \in [0, 1]$,

$$\bar{m}(\alpha \mathbf{x} + (1 - \alpha) \mathbf{x}', \mathbf{y}) \subseteq \alpha \bar{m}(\mathbf{x}, \mathbf{y}) + (1 - \alpha) \bar{m}(\mathbf{x}', \mathbf{y}).$$

A function $\bar{m} : \Delta(\mathcal{A}) \times \Delta(\mathcal{B}) \rightarrow \mathcal{S}(\mathbb{R}^d)$ is convex in its second argument if for all $\mathbf{x} \in \Delta(\mathcal{A})$, all $\mathbf{y}, \mathbf{y}' \in \Delta(\mathcal{B})$, and all $\alpha \in [0, 1]$,

$$\alpha \bar{m}(\mathbf{x}, \mathbf{y}) + (1 - \alpha) \bar{m}(\mathbf{x}, \mathbf{y}') \subseteq \bar{m}(\mathbf{x}, \alpha \mathbf{y} + (1 - \alpha) \mathbf{y}').$$

An example of such a concave-convex function \bar{m} is discussed in Lemma 17.

The following theorem indicates that (SVAC) is the necessary and sufficient condition for the \bar{m} -approachability of a closed convex set \mathcal{C} when the payoff function \bar{m} satisfies all four properties stated in Definitions 13 and 14. (Boundedness of \bar{m} indeed follows from the continuity of \bar{m} in each variable.)

Theorem 15 If \bar{m} is bounded, convex, and uniformly continuous in its second argument, then (SVAC) entails that a closed convex set \mathcal{C} is \bar{m} -approachable.

If \bar{m} is concave and uniformly continuous in its first argument, then a closed convex set \mathcal{C} is \bar{m} -approachable only if (SVAC) is satisfied.

The proof of the necessity statement follows closely the arguments used in the proof of Theorem 8. The sufficiency statement relies on the use of what is called a calibrated strategy, where we define calibration in a (slightly) stronger way than Foster and Vohra (1998) did. All the details, including the definition of the stronger notion of calibration and the construction of an algorithm controlling it, can be found in Appendix A.

5. Approachability in Games with Partial Monitoring

A repeated vector-valued game with partial monitoring is described as follows (see, e.g., Mertens et al., 1994, Rustichini, 1999, and references therein). The players have respective finite action sets \mathcal{I} and \mathcal{J} . We denote by $r : \mathcal{I} \times \mathcal{J} \rightarrow \mathbb{R}^d$ the vector-valued payoff function of the first player and extend it multi-linearly to $\Delta(\mathcal{I}) \times \Delta(\mathcal{J})$. At each round, players simultaneously choose their actions $I_t \in \mathcal{I}$ and $J_t \in \mathcal{J}$, possibly at random according to probability distributions denoted by $\mathbf{p}_t \in \Delta(\mathcal{I})$ and $\mathbf{q}_t \in \Delta(\mathcal{J})$. At the end of a round, the

first player does not observe J_t nor $r(I_t, J_t)$ but only receives a signal. There is a finite set \mathcal{H} of possible signals; the feedback S_t that is given to the first player is drawn at random according to the distribution $H(I_t, J_t)$, where the mapping $H : \mathcal{I} \times \mathcal{J} \rightarrow \Delta(\mathcal{H})$ is known by the first player. We will refer to H as the signaling structure.

Formally, strategies of the first player now associate with I_1, \dots, I_{t-1} and S_1, \dots, S_{t-1} a mixed action $\mathbf{p}_t \in \Delta(\mathcal{I})$, according to which I_t is drawn independently. We do not impose any restriction on the opponent player, who enjoys a full monitoring: strategies of his associate with I_1, \dots, I_{t-1} , with J_1, \dots, J_{t-1} and with S_1, \dots, S_{t-1} a mixed action $\mathbf{q}_t \in \Delta(\mathcal{I})$, according to which J_t is drawn independently.

Example 1 *Examples of such partial monitoring games are provided by, e.g., Cesa-Bianchi et al. (2006), among which we can cite the apple tasting problem, the label-efficient prediction constraint, and the multi-armed bandit settings.*

Some additional notation will be useful. We denote by R a bound on the norm of (the linear extension of) r ,

$$R = \max_{(i,j) \in \mathcal{I} \times \mathcal{J}} \|r(i, j)\|_2. \tag{5}$$

The cardinalities of the finite sets \mathcal{I} , \mathcal{J} , and \mathcal{H} will be referred to as $N_{\mathcal{I}}$, $N_{\mathcal{J}}$, and $N_{\mathcal{H}}$.

The definition of approachability can be extended from the setting of full information to the setting of partial monitoring as follows. The only new ingredient is the signaling structure H , the aim is unchanged.

Definition 16 *Let $\mathcal{C} \subseteq \mathbb{R}^d$ be some set; \mathcal{C} is r -approachable by the first player for the signaling structure H if he has a strategy such that, for all $\varepsilon > 0$, there exists an integer T_ε such that for all strategies of the second player,*

$$\mathbb{P} \left\{ \forall T \geq T_\varepsilon, \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2 \leq \varepsilon \right\} \geq 1 - \varepsilon.$$

In particular, the first player has a strategy that ensures that the sequence of his average vector-valued payoffs converges almost surely to the set \mathcal{C} (uniformly with respect to the strategies of the second player), even if he only observes the random signals S_t as a feedback.

Here again, more precise approachability guarantees than the ones required by the definition will be obtained. Indeed, Corollary 27 exhibits bounds of the following form, for a suitable strategy of the first player. For all strategies of the second player and for all $T \geq 1$, with probability at least $1 - P_T$,

$$\sup_{\tau \geq T} \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{\tau} \sum_{t=1}^{\tau} r(I_t, J_t) \right\|_2 \leq R_T,$$

where $P_T = O(1/T)$ and $R_T = O(T^{-1/5} \ln T)$.

Our contributions to approachability in games with partial monitoring: A necessary and sufficient condition for r -approachability with the signaling structure H was already stated

and proved by Perchet (2011a), together with an approaching strategy. We therefore need to detail where our contribution lies.

First, our strategy is efficient (as long as some projection operator can be computed efficiently, e.g., in the case when the target set is a polytope, see Sections 5.4.1–5.4.2 as well as in the cases of external and internal regret minimization described below in Section 6). In contrast, the one of Perchet (2011a) relies on auxiliary strategies that are calibrated and that require a grid that is progressively refined (leading to a step complexity that is prohibitive in the number T of past steps and to rates of convergence that become dependent on the dimension). The latter construction is in essence the one used in Section 4.

Second, we are able, for the first time, to exhibit convergence rates that are independent of the dimension (as is the case with full monitoring). A somehow related result appeared in Perchet (2011b), but only for the special case of regret minimization. The proof techniques used therein are involved and hold only for regret minimization, not for general approachability.

Third, as far as elegance is concerned, our proof of the sufficiency of the condition for r -approachability with the signaling structure H is short, compact, and more direct than the one of Perchet (2011a) or even of Perchet (2011b), which relied on several layers of concepts (for example, calibration or internal regret in games with partial monitoring).

5.1 Statement of the Necessary and Sufficient Condition for Approachability

To recall the mentioned approachability condition of Perchet (2011a) we need some additional notation. For all $\mathbf{q} \in \Delta(\mathcal{J})$, we denote by $\tilde{H}(\mathbf{q})$ the element in $\Delta(\mathcal{H})^{\mathcal{I}}$ defined as follows. For all $i \in \mathcal{I}$, its i -th component is given by the convex combination of probability distributions over \mathcal{H}

$$\tilde{H}(\mathbf{q})_i = H(i, \mathbf{q}) = \sum_{j \in \mathcal{J}} q_j H(i, j).$$

Also, we denote by \mathcal{F} the convex set of feasible vectors of probability distributions over \mathcal{H} :

$$\mathcal{F} = \left\{ \tilde{H}(\mathbf{q}) : \mathbf{q} \in \Delta(\mathcal{J}) \right\}.$$

A generic element of \mathcal{F} will be denoted by $\sigma \in \mathcal{F}$ and we define the set-valued function \bar{m} , for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$, by

$$\bar{m}(\mathbf{p}, \sigma) = \left\{ r(\mathbf{p}, \mathbf{q}_{\text{eqv}}) : \mathbf{q}_{\text{eqv}} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}_{\text{eqv}}) = \sigma \right\}.$$

We use in \mathbf{q}_{eqv} the subscript “eqv” (standing for “equivalent”) as all considered \mathbf{q}_{eqv} vectors induce the same distributions of signals σ and are thus equivalent from the monitoring perspective.

The necessary and sufficient condition exhibited by Perchet (2011a) for the r -approachability of \mathcal{C} with the signaling structure H can now be recalled. In the sequel we will refer to this condition as Condition (APM), an acronym that stands for “approachability with partial monitoring.”

Condition 1 [referred to as Condition (APM)] *The signaling structure H , the vector-payoff function r , and the set \mathcal{C} satisfy*

$$\forall \mathbf{q} \in \Delta(\mathcal{J}), \exists \mathbf{p} \in \Delta(\mathcal{I}), \forall \mathbf{q}' \in \Delta(\mathcal{J}), \quad \tilde{H}(\mathbf{q}) = \tilde{H}(\mathbf{q}') \Rightarrow r(\mathbf{p}, \mathbf{q}') \in \mathcal{C}.$$

The condition can be equivalently reformulated as

$$\forall \sigma \in \mathcal{F}, \quad \exists \mathbf{p} \in \Delta(\mathcal{I}), \quad \bar{m}(\mathbf{p}, \sigma) \subseteq \mathcal{C}. \quad (\text{APM})$$

The subsequent sections show (in a constructive way, with a strategy efficient up to a projection oracle) that Condition (APM) is sufficient for r -approachability of closed convex sets \mathcal{C} given the signaling structure H . That this condition is necessary was already proved in Section 3.1 of Perchet (2011a).

5.2 Links with Set-Valued Approachability

As will become clear in the proof of Theorem 24, the key in our problem will be to ensure the set-valued approachability of \mathcal{C} with the following non-linear set-valued payoff function, that is however concave-convex in the sense of Definition 14.

Lemma 17 *The function*

$$(\mathbf{p}, \mathbf{q}) \in \Delta(\mathcal{I}) \times \Delta(\mathcal{J}) \longmapsto \bar{m}(\mathbf{p}, \tilde{H}(\mathbf{q}))$$

is concave in its first argument and convex in its second argument.

Proof For the concavity part, we consider some pair $\mathbf{p}, \mathbf{p}' \in \Delta(\mathcal{I})$, some $\mathbf{q} \in \Delta(\mathcal{J})$ and some $\alpha \in [0, 1]$. By the linearity of r , the elements of the set of interest can be written as

$$\begin{aligned} & \bar{m}(\alpha \mathbf{p} + (1 - \alpha) \mathbf{p}', \tilde{H}(\mathbf{q})) \\ &= \left\{ \alpha r(\mathbf{p}, \mathbf{q}_{\text{eqv}}) + (1 - \alpha) r(\mathbf{p}', \mathbf{q}_{\text{eqv}}) : \mathbf{q}_{\text{eqv}} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}_{\text{eqv}}) = \tilde{H}(\mathbf{q}) \right\}. \end{aligned}$$

This set is therefore indeed included in (but in general, not equal to)

$$\begin{aligned} & \alpha \bar{m}(\mathbf{p}, \tilde{H}(\mathbf{q})) + (1 - \alpha) \alpha \bar{m}(\mathbf{p}, \tilde{H}(\mathbf{q})) \\ &= \alpha \left\{ r(\mathbf{p}, \mathbf{q}_{\text{eqv}}) : \mathbf{q}_{\text{eqv}} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}_{\text{eqv}}) = \tilde{H}(\mathbf{q}) \right\} \\ & \quad + (1 - \alpha) \left\{ r(\mathbf{p}, \mathbf{q}'_{\text{eqv}}) : \mathbf{q}'_{\text{eqv}} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}'_{\text{eqv}}) = \tilde{H}(\mathbf{q}) \right\}. \end{aligned}$$

Similarly, for the convexity part, we consider some pair $\mathbf{q}, \mathbf{q}' \in \Delta(\mathcal{J})$, some $\mathbf{p} \in \Delta(\mathcal{I})$ and some $\alpha \in [0, 1]$. Elements of the convex combination of sets

$$\alpha \bar{m}(\mathbf{p}, \tilde{H}(\mathbf{q})) + (1 - \alpha) \bar{m}(\mathbf{p}, \tilde{H}(\mathbf{q}'))$$

are of the form

$$\alpha r(\mathbf{p}, \mathbf{q}_{\text{eqv}}) + (1 - \alpha) r(\mathbf{p}, \mathbf{q}'_{\text{eqv}}) = r(\mathbf{p}, \alpha \mathbf{q}_{\text{eqv}} + (1 - \alpha) \mathbf{q}'_{\text{eqv}}),$$

where \mathbf{q}_{eqv} and \mathbf{q}'_{eqv} are such that

$$\tilde{H}(\mathbf{q}_{\text{eqv}}) = \tilde{H}(\mathbf{q}) \quad \text{and} \quad \tilde{H}(\mathbf{q}'_{\text{eqv}}) = \tilde{H}(\mathbf{q}'). \quad (6)$$

In particular, by linearity of \tilde{H} , we have

$$\tilde{H}(\alpha \mathbf{q}_{\text{eqv}} + (1 - \alpha) \mathbf{q}'_{\text{eqv}}) = \tilde{H}(\alpha \mathbf{q} + (1 - \alpha) \mathbf{q}'),$$

which shows that

$$r(\mathbf{p}, \alpha \mathbf{q}_{\text{eqv}} + (1 - \alpha) \mathbf{q}'_{\text{eqv}}) \in \bar{m}(\mathbf{p}, \tilde{H}(\alpha \mathbf{q} + (1 - \alpha) \mathbf{q}')).$$

The desired inclusion

$$\alpha \bar{m}(\mathbf{p}, \tilde{H}(\mathbf{q})) + (1 - \alpha) \bar{m}(\mathbf{p}, \tilde{H}(\mathbf{q}')) \subseteq \bar{m}(\mathbf{p}, \tilde{H}(\alpha \mathbf{q} + (1 - \alpha) \mathbf{q}'))$$

follows. Note that this inclusion is not an equality in general, as it cannot be guaranteed that any $\mathbf{q}''_{\text{eqv}}$ such that

$$\tilde{H}(\alpha \mathbf{q} + (1 - \alpha) \mathbf{q}') = \tilde{H}(\mathbf{q}''_{\text{eqv}})$$

can be decomposed under the form $\alpha \mathbf{q}_{\text{eqv}} + (1 - \alpha) \mathbf{q}'_{\text{eqv}}$, where \mathbf{q}_{eqv} and \mathbf{q}'_{eqv} satisfy (6). ■

Unfortunately, efficient strategies for set-valued approachability were only proposed in the linear case (Section 3), not in the concave–convex case (Section 4), and the proof of Lemma 17 shows that linearity cannot be guaranteed per se. However, we illustrate in the next example (and provide a general theory in the next section) how working in lifted spaces can lead to linearity and hence to efficiency.

Example 2 *We consider a game in which the second player (the column player) can force the first player (the row player) to play a game of matching pennies in the dark by choosing actions L or M . More formally, in the matrix below, the real numbers denote the payoff while \clubsuit and \heartsuit denote the two possible signals. The respective sets of actions are $\mathcal{I} = \{T, B\}$ and $\mathcal{J} = \{L, M, R\}$.*

	L	M	R
T	$1 / \clubsuit$	$-1 / \clubsuit$	$2 / \heartsuit$
B	$-1 / \clubsuit$	$1 / \clubsuit$	$3 / \heartsuit$

In this example we only study the mapping $\mathbf{p} \mapsto \bar{m}(\mathbf{p}, \clubsuit)$ and show that it is piecewise linear on $\Delta(\mathcal{I})$, thus, is induced by a linear mapping defined on a lifted space.

We introduce a set $\mathcal{A} = \{\mathbf{p}_T, \mathbf{p}_B, \mathbf{p}_{1/2}\}$ of possibly mixed actions extending the set $\mathcal{I} = \{T, B\}$ of pure actions; the set \mathcal{A} is composed of

$$\mathbf{p}_T = \delta_T, \quad \mathbf{p}_B = \delta_B, \quad \text{and} \quad \mathbf{p}_{1/2} = \frac{1}{2} \delta_T + \frac{1}{2} \delta_B.$$

Each mixed action in $\Delta(\mathcal{I})$ can be uniquely written as $\mathbf{p}_\lambda = \lambda \delta_B + (1 - \lambda) \delta_T$ for some $\lambda \in [0, 1]$. Now, for $\lambda \geq 1/2$, first,

$$\mathbf{p}_\lambda = (2\lambda - 1) \delta_B + (1 - (2\lambda - 1)) \mathbf{p}_{1/2};$$

second, by definition of \bar{m} ,

$$\bar{m}(\mathbf{p}_\lambda, \clubsuit) = [1 - 2\lambda, 2\lambda - 1];$$

since in particular $\bar{m}(\mathbf{p}_{1/2}, \clubsuit) = \{0\}$ and $\bar{m}(\delta_B, \clubsuit) = [-1, 1]$, we have the convex decomposition

$$\bar{m}(\mathbf{p}_\lambda, \clubsuit) = (2\lambda - 1)\bar{m}(\delta_B, \clubsuit) + (1 - (2\lambda - 1))\bar{m}(\mathbf{p}_{1/2}, \clubsuit),$$

that can be restated as

$$\bar{m}\left((2\lambda - 1)\delta_B + (1 - (2\lambda - 1))\mathbf{p}_{1/2}, \clubsuit\right) = (2\lambda - 1)\bar{m}(\delta_B, \clubsuit) + (1 - (2\lambda - 1))\bar{m}(\mathbf{p}_{1/2}, \clubsuit).$$

That is, $\bar{m}(\cdot, \clubsuit)$ is linear on the subset of $\Delta(\mathcal{I})$ corresponding to mixed actions \mathbf{p}_λ with $\lambda \geq 1/2$.

A similar property holds on the subset of distributions with $\lambda \leq 1/2$, so that we proved that $\bar{m}(\cdot, \clubsuit)$ is piecewise linear on $\Delta(\mathcal{I})$.

The linearity on a lifted space comes from the following observation: \bar{m} is induced by the linear extension to $\Delta(\mathcal{A})$ of the restriction of \bar{m} to \mathcal{A} (see Definition 21 for a more formal statement).

5.3 A Particular Class of Games, Encompassing Regret Minimization

In this section we consider the case where the payoff function and the signaling structure have some special properties described below (linked to linearity properties on lifted spaces and called “bi-piecewise linearity”) and that can be exploited to get efficient strategies. The case of general games with partial monitoring is then considered in Section 5.4 but the particular class of games considered here is already rich enough to encompass the minimization of external and internal regret, as will be seen in Section 6.

To define bi-piecewise linearity of a game with partial monitoring, we start from a technical lemma that shows that $\bar{m}(\mathbf{p}, \sigma)$ can be written as a *finite* convex combination of sets of the form $\bar{m}(\mathbf{p}, b)$, where b belongs to some finite set $\mathcal{B} \subseteq \mathcal{F}$ that depends on the game. Under the additional assumption of the so-called piecewise linearity of the thus-defined mappings $\bar{m}(\cdot, b)$, we then describe an efficient strategy for approachability (up to a projection oracle) followed by convergence rate guarantees.

Definition 18 *Let P be a polytope and let \mathcal{X} be a convex set. A mapping $f : P \rightarrow \mathcal{X}$ is piecewise linear if f is continuous and*

- *there exist finitely many sub-polytopes P_1, \dots, P_K covering P and such that two different sub-polytopes $P_k, P_{k'}$ have an intersection with empty interior; we call these sub-polytopes a decomposition of P ;*
- *f is linear on each sub-polytope P_k .*

5.3.1 BI-PIECEWISE LINEARITY OF A GAME — A PRELIMINARY TECHNICAL RESULT

Lemma 19 *For any game with partial monitoring, there exists a finite set $\mathcal{B} \subset \mathcal{F}$ and a piecewise-linear (injective) mapping $\Phi : \mathcal{F} \rightarrow \Delta(\mathcal{B})$ such that*

$$\forall \sigma \in \mathcal{F}, \quad \forall \mathbf{p} \in \Delta(\mathcal{I}), \quad \bar{m}(\mathbf{p}, \sigma) = \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \bar{m}(\mathbf{p}, b),$$

where we denoted the convex weight vector $\Phi(\sigma) \in \Delta(\mathcal{B})$ by $(\Phi_b(\sigma))_{b \in \mathcal{B}}$.

Proof \tilde{H} is linear on the polytope $\Delta(\mathcal{J})$; Proposition 2.4 in Rambau and Ziegler (1996) thus implies that its inverse mapping \tilde{H}^{-1} is a piecewise linear mapping of \mathcal{F} into the set of the subsets of $\Delta(\mathcal{J})$. (Note that the latter set has a structure of a convex set, see Footnote 1.) This means by definition that there exists a finite decomposition of \mathcal{F} into polytopes P_1, \dots, P_K each on which \tilde{H}^{-1} is linear. Up to a triangulation (see, e.g., Goodman and O'Rourke, 2004, Chapter 14), we can assume that each P_k is a simplex. Denote by $\mathcal{B}_k \subseteq \mathcal{F}$ the set of vertices of P_k ; then, the finite subset stated in the lemma is

$$\mathcal{B} = \bigcup_{k=1}^K \mathcal{B}_k,$$

the set of all vertices of all the simplices.

Fix any $\sigma \in \mathcal{F}$. It belongs to some simplex P_k , so that there exists a convex decomposition $\sigma = \sum_{b \in \mathcal{B}_k} \lambda_b b$; this decomposition is unique within the simplex P_k . If σ belongs to two different simplices, then it actually belongs to their common face and the two possible decompositions coincide (some coefficients λ_b in the above decomposition are null). All in all, with each $\sigma \in \mathcal{F}$, we can associate a unique decomposition in \mathcal{B} ,

$$\sigma = \sum_{b \in \mathcal{B}} \Phi_b(\sigma) b,$$

where the coefficients $(\Phi_b(\sigma))_{b \in \mathcal{B}}$ form a convex weight vector over \mathcal{B} , i.e., belong to $\Delta(\mathcal{B})$; in addition, $\Phi_b(\sigma) > 0$ only if $b \in \mathcal{B}_k$, where k is such that $\sigma \in P_k$.

Since \tilde{H}^{-1} is linear on each simplex P_1, \dots, P_K , we therefore get

$$\tilde{H}^{-1}(\sigma) = \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \tilde{H}^{-1}(b).$$

Finally, the result is a consequence of the fact that

$$\bar{m}(\mathbf{p}, \sigma) = r\left(\mathbf{p}, \tilde{H}^{-1}(\sigma)\right) = r\left(\mathbf{p}, \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \tilde{H}^{-1}(b)\right),$$

that implies, by the linearity of r , that

$$\bar{m}(\mathbf{p}, \sigma) = \sum_{b \in \mathcal{B}} \Phi_b(\sigma) r\left(\mathbf{p}, \tilde{H}^{-1}(b)\right) = \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \bar{m}(\mathbf{p}, b),$$

which concludes the proof. ■

Remark 20 *The proof shows that Φ is piecewise linear on a finite decomposition of \mathcal{F} ; it is therefore Lipschitz on \mathcal{F} . We denote by κ_Φ its Lipschitz constant with respect to the ℓ^2 -norms.*

The main contribution of this subsection (Definition 21) relies on the following additional assumption.

Assumption 1 *We assume that $\bar{m}(\cdot, b)$ is piecewise linear on $\Delta(\mathcal{I})$ for every $b \in \mathcal{B}$. We then call the corresponding game (r, H) a bi-piecewise linear game.*

Assumption 1 means that for each $b \in \mathcal{B}$ there exists a decomposition of $\Delta(\mathcal{I})$ into polytopes each on which $\bar{m}(\cdot, b)$ is linear. Since \mathcal{B} is finite, there exist finitely many such decompositions to consider, and thus there exists a decomposition to polytopes that refines all of them. (The latter is generated by the intersection of all considered polytopes as b varies.) By construction, every $\bar{m}(\cdot, b)$ is linear on any of the polytopes of this common decomposition. We denote by $\mathcal{A} \subset \Delta(\mathcal{I})$ the finite subset of all their vertices. A construction similar to the one used in the proof of Lemma 19 leads to a piecewise linear (injective) mapping $\Theta : \Delta(\mathcal{I}) \rightarrow \Delta(\mathcal{A})$, where $\Theta(\mathbf{p})$ is the decomposition of \mathbf{p} on the vertices of the polytope(s) of the decomposition to which it belongs, satisfying

$$\forall b \in \mathcal{B}, \quad \forall \mathbf{p} \in \Delta(\mathcal{I}), \quad \bar{m}(\mathbf{p}, b) = \sum_{a \in \mathcal{A}} \Theta_a(\mathbf{p}) \bar{m}(a, b),$$

where we denoted the convex weight vector $\Theta(\mathbf{p}) \in \Delta(\mathcal{A})$ by $(\Theta_a(\mathbf{p}))_{a \in \mathcal{A}}$. This, Lemma 19, and Assumption 1 show that on a lifted space, \bar{m} coincides with a bi-linear mapping $\overline{\bar{m}}$, as is made formal in the next definition.

Definition 21 *For a bi-piecewise linear game, we denote by $\overline{\bar{m}}$ the linear extension to $\Delta(\mathcal{A} \times \mathcal{B})$ of the restriction of \bar{m} to $\mathcal{A} \times \mathcal{B}$, so that for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$,*

$$\overline{\bar{m}}(\mathbf{p}, \sigma) = \overline{\bar{m}}(\Theta(\mathbf{p}), \Phi(\sigma)).$$

5.3.2 CONSTRUCTION OF A STRATEGY TO APPROACH \mathcal{C}

The approaching strategy for the original problem is based on a strategy Ψ for $\overline{\bar{m}}$ -approachability of \mathcal{C} , provided by Theorem 8; we therefore first need to prove the existence of such a Ψ .

Lemma 22 *Under Condition (APM), the closed convex set \mathcal{C} is $\overline{\bar{m}}$ -approachable.*

Proof We show that Condition (SVAC) in Theorem 8 is satisfied, that is, that for all $\mathbf{y} \in \Delta(\mathcal{B})$, there exists some $\mathbf{x} \in \Delta(\mathcal{A})$ such that $\overline{\bar{m}}(\mathbf{x}, \mathbf{y}) \subseteq \mathcal{C}$. With such a given $\mathbf{y} \in \Delta(\mathcal{B})$, we associate² the feasible vector of signals $\sigma = \sum_{b \in \mathcal{B}} y_b b \in \mathcal{F}$ and let \mathbf{p} be given by Condition (APM), so that $\overline{\bar{m}}(\mathbf{p}, \sigma) \subseteq \mathcal{C}$. By linearity of $\overline{\bar{m}}$ (for the first equality), by

2. Note, however, that we do not necessarily have that $\Phi(\sigma)$ and \mathbf{y} are equal, as Φ is not a one-to-one mapping (it is injective but not surjective).

 Approaching Strategy in Games with Partial Monitoring

Parameters: an integer block length $L \geq 1$, an exploration parameter $\gamma \in [0, 1]$, a strategy Ψ for \overline{m} -approachability of \mathcal{C}

Notation: $\mathbf{u} \in \Delta(\mathcal{I})$ is the uniform distribution over \mathcal{I} , $P_{\mathcal{F}}$ denotes the projection operator in ℓ^2 -norm of $\mathbb{R}^{\mathcal{H} \times \mathcal{I}}$ onto \mathcal{F}

Initialization: compute the finite set \mathcal{B} and the mapping $\Phi : \mathcal{F} \rightarrow \Delta(\mathcal{B})$ of Lemma 19, compute the finite set \mathcal{A} and the mapping $\Theta : \Delta(\mathcal{I}) \rightarrow \Delta(\mathcal{A})$ defined based on Assumption 1, pick an arbitrary $\theta_1 \in \Delta(\mathcal{A})$

For all blocks $n = 1, 2, \dots$,

1. define $\mathbf{x}_n = \sum_{a \in \mathcal{A}} \theta_{n,a} a$ and $\mathbf{p}_n = (1 - \gamma) \mathbf{x}_n + \gamma \mathbf{u}$; refer to the components of \mathbf{p}_n as $(p_{i,n})_{i \in \mathcal{I}}$;
2. for rounds $t = (n - 1)L + 1, \dots, nL$,
 - 2.1 draw an action $I_t \in \mathcal{I}$ at random according to \mathbf{p}_n ;
 - 2.2 get the signal S_t ;
3. form the estimated vector of probability distributions over signals,

$$\tilde{\sigma}_n = \left(\frac{1}{L} \sum_{t=(n-1)L+1}^{nL} \frac{\mathbb{I}_{\{S_t=s\}} \mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} \right)_{(i,s) \in \mathcal{I} \times \mathcal{H}} ;$$

4. compute the projection $\hat{\sigma}_n = P_{\mathcal{F}}(\tilde{\sigma}_n)$;
 5. choose $\theta_{n+1} = \Psi(\theta_1, \Phi(\hat{\sigma}_1), \dots, \theta_n, \Phi(\hat{\sigma}_n))$.
-

Figure 1: The proposed strategy, which plays in blocks.

convexity of \overline{m} in its second argument (for the first inclusion), by Lemma 19 (for the second and fourth equalities), by construction of \mathcal{A} (for the third equality),

$$\begin{aligned} \overline{m}(\Theta(\mathbf{p}), \mathbf{y}) &= \sum_{a \in \mathcal{A}} \Theta_a(\mathbf{p}) \sum_{b \in \mathcal{B}} y_b \overline{m}(a, b) \subseteq \sum_{a \in \mathcal{A}} \Theta_a(\mathbf{p}) \overline{m}(a, \sigma) = \sum_{a \in \mathcal{A}} \Theta_a(\mathbf{p}) \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \overline{m}(a, b) \\ &= \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \overline{m}(\mathbf{p}, b) = \overline{m}(\mathbf{p}, \sigma) \subseteq \mathcal{C}, \end{aligned}$$

which concludes the proof. ■

We consider the strategy described in Figure 1 (and the notation introduced therein). It forces exploration at a γ rate, as is usual in situations with partial monitoring. One of its key ingredients, that conditionally unbiased estimators are available, is extracted from Lugosi et al. (2008, Section 6): in block n we consider sums of elements of the form

$$\hat{H}_t = \left(\frac{\mathbb{I}_{\{S_t=s\}} \mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} \right)_{(i,s) \in \mathcal{I} \times \mathcal{H}} \in \mathbb{R}^{\mathcal{H} \times \mathcal{I}}.$$

Averaging over the respective random draws of I_t and S_t according to \mathbf{p}_n and $H(I_t, J_t)$, i.e., taking the conditional expectation \mathbb{E}_t with respect to \mathbf{p}_n and J_t , we get

$$\mathbb{E}_t[\widehat{H}_t] = \widetilde{H}(\delta_{J_t}). \quad (7)$$

Indeed, the conditional expectation of the component i of \widehat{H}_t equals

$$\mathbb{E}_t \left[\left(\frac{\mathbb{I}_{\{S_t=s\}} \mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} \right)_{s \in \mathcal{H}} \right] = \mathbb{E}_t \left[\frac{H(I_t, J_t) \mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} \right] = \frac{H(i, J_t)}{p_{i,n}} \mathbb{E}_t[\mathbb{I}_{\{I_t=i\}}] = H(i, J_t),$$

where we first took the expectation over the random draw of S_t (conditionally on \mathbf{p}_n , J_t , and I_t) and then over the one of I_t . Consequently, concentration arguments show that for L large enough,

$$\widetilde{\sigma}_n = \frac{1}{L} \sum_{t=(n-1)L+1}^{nL} \widehat{H}_t \quad \text{is close to} \quad \widetilde{H}(\widehat{\mathbf{q}}_n), \quad \text{where} \quad \widehat{\mathbf{q}}_n = \frac{1}{L} \sum_{t=(n-1)L+1}^{nL} \delta_{J_t}. \quad (8)$$

Actually, since $\mathcal{F} \subseteq \Delta(\mathcal{H})^{\mathcal{I}}$, we have a natural embedding of \mathcal{F} into $\mathbb{R}^{\mathcal{H} \times \mathcal{I}}$ and we can define $P_{\mathcal{F}}$, the convex projection operator onto \mathcal{F} (in ℓ^2 -norm). Instead of using directly $\widetilde{\sigma}_n$, we consider in our strategy $\widehat{\sigma}_n = P_{\mathcal{F}}(\widetilde{\sigma}_n)$, which is even closer to $\widetilde{H}(\widehat{\mathbf{q}}_n)$.

More precisely, the following result can be extracted from the proof of Theorem 6.1 in Lugosi et al. (2008). For the convenience of the reader, a self-contained proof is provided in Appendix C.

Lemma 23 *For all $\delta \in (0, 1)$, for all blocks $n \geq 1$, with probability at least $1 - \delta$,*

$$\left\| \widehat{\sigma}_n - \widetilde{H}(\widehat{\mathbf{q}}_n) \right\|_2 \leq \sqrt{N_{\mathcal{I}} N_{\mathcal{H}}} \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta} \right).$$

5.3.3 A PERFORMANCE GUARANTEE FOR THE STRATEGY OF FIGURE 1

For the sake of simplicity, we provide first a performance bound for fixed parameters γ and L tuned as functions of a known horizon T . We then obtain a bound holding only for the specific round T . Adaptation to $T \rightarrow \infty$ (and the obtention of bounds for all $T \geq 1$) are then described in the next section. We recall that R was defined in (5) as a bound in norm on r .

Theorem 24 *Consider a closed convex set \mathcal{C} and a game (r, H) for which Condition (APM) is satisfied and that is bi-piecewise linear in the sense of Assumption 1. Then, for all strategies of the second player, the strategy of Figure 1, run with parameters $\gamma \in [0, 1]$ and $L \geq 1$ and fed with a strategy Ψ for \overline{m} -approachability of \mathcal{C} (provided by Lemma 22) is such that, for all $T \geq L + 1$, for all $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2 \leq \frac{2L}{T} R + 4R \sqrt{\frac{\ln(2(T+L)/(L\delta))}{T-L}} + 2\gamma R + \frac{2R}{\sqrt{T/L-1}} \sqrt{N_{\mathcal{A}} N_{\mathcal{B}}} \\ + R \kappa_{\Phi} \sqrt{N_{\mathcal{I}} N_{\mathcal{H}} N_{\mathcal{A}}} \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}} (T+L)}{L\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}} (T+L)}{L\delta} \right).$$

In particular, for all $T \geq 1$, the choices of $L = \lceil T^{3/5} \rceil$ and $\gamma = T^{-1/5}$ imply that for all strategies of the second player, for all $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2 \leq \Xi \left(T^{-1/5} \sqrt{\ln \frac{T}{\delta}} + T^{-2/5} \ln \frac{T}{\delta} \right)$$

for some constant Ξ depending only on \mathcal{C} and on the game (r, H) at hand.

The efficiency of the strategy of Figure 1 depends on whether it can be fed with an efficient approaching strategy Ψ , which in turn depends on the respective geometries of \bar{m} and \mathcal{C} , as was indicated at the end of Section 3.1. (Note that the projection onto \mathcal{F} can be performed in polynomial time, as the latter closed convex set is defined by finitely many linear constraints, and that the computation of \mathcal{A} , \mathcal{B} , and \bar{m} can be performed beforehand.) In any case, the per-round complexity is constant (though possibly large).

Proof We write T as $T = NL + k$ where N is an integer and $0 \leq k \leq L - 1$ and will show successively that (possibly with overwhelming probability only) the following statements hold.

$$\frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \quad \text{is close to} \quad \frac{1}{NL} \sum_{t=1}^{NL} r(I_t, J_t); \tag{9}$$

$$\frac{1}{NL} \sum_{t=1}^{NL} r(I_t, J_t) \quad \text{is close to} \quad \frac{1}{N} \sum_{n=1}^N r(\mathbf{p}_n, \hat{\mathbf{q}}_n); \tag{10}$$

$$\frac{1}{N} \sum_{n=1}^N r(\mathbf{p}_n, \hat{\mathbf{q}}_n) \quad \text{is close to} \quad \frac{1}{N} \sum_{n=1}^N r(\mathbf{x}_n, \hat{\mathbf{q}}_n); \tag{11}$$

$$\begin{aligned} & \frac{1}{N} \sum_{n=1}^N r(\mathbf{x}_n, \hat{\mathbf{q}}_n) \\ &= \frac{1}{N} \sum_{n=1}^N \sum_{a \in \mathcal{A}} \theta_{n,a} r(a, \hat{\mathbf{q}}_n) \quad \text{belongs to the set} \quad \frac{1}{N} \sum_{n=1}^N \sum_{a \in \mathcal{A}} \theta_{n,a} \bar{m}(a, \tilde{H}(\hat{\mathbf{q}}_n)); \\ & \frac{1}{N} \sum_{n=1}^N \sum_{a \in \mathcal{A}} \theta_{n,a} \bar{m}(a, \tilde{H}(\hat{\mathbf{q}}_n)) \quad \text{is equal to the set} \quad \frac{1}{N} \sum_{n=1}^N \bar{m}(\boldsymbol{\theta}_n, \Phi(\tilde{H}(\hat{\mathbf{q}}_n))); \\ & \frac{1}{N} \sum_{n=1}^N \bar{m}(\boldsymbol{\theta}_n, \Phi(\tilde{H}(\hat{\mathbf{q}}_n))) \quad \text{is close to the set} \quad \frac{1}{N} \sum_{n=1}^N \bar{m}(\boldsymbol{\theta}_n, \Phi(\hat{\sigma}_n)); \tag{12} \end{aligned}$$

$$\frac{1}{N} \sum_{n=1}^N \bar{m}(\boldsymbol{\theta}_n, \Phi(\hat{\sigma}_n)) \quad \text{is close to the set} \quad \mathcal{C}; \tag{13}$$

where we recall that the notation $\hat{\mathbf{q}}_n$ was defined in (8) and is referring to the empirical distribution of the J_t in the n -th block. Actually, we will show below the numbered statements only. The first unnumbered statement is immediate by the definition of \mathbf{x}_n , the linearity of

r , and the very definition of \bar{m} ; while the second one follows from Definition 21:

$$\begin{aligned} \frac{1}{N} \sum_{n=1}^N \sum_{a \in \mathcal{A}} \theta_{n,a} \bar{m}(a, \tilde{H}(\hat{\mathbf{q}}_n)) &= \frac{1}{N} \sum_{n=1}^N \sum_{(a,b) \in \mathcal{A} \times \mathcal{B}} \theta_{n,a} \Phi_b(\tilde{H}(\hat{\mathbf{q}}_n)) \bar{m}(a,b) \\ &= \frac{1}{N} \sum_{n=1}^N \bar{m}(\boldsymbol{\theta}_n, \Phi(\tilde{H}(\hat{\mathbf{q}}_n))). \end{aligned}$$

Step 1: Assertion (9). A direct calculation decomposing the sum over T elements into a sum over the NL first elements and the k remaining ones shows that

$$\left\| \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) - \frac{1}{NL} \sum_{t=1}^{NL} r(I_t, J_t) \right\|_2 \leq R \left(\frac{k}{T} + \left(\frac{1}{NL} - \frac{1}{T} \right) NL \right) = \frac{2k}{T} R \leq \frac{2L}{T} R.$$

Step 2: Assertion (10). We note that by defining \mathbb{E}_t as the conditional expectation with respect to $(I_1, S_1, J_1), \dots, (I_{t-1}, S_{t-1}, J_{t-1})$ and J_t , which fixes the values of the distribution \mathbf{p}'_t of I_t and the value of J_t , we have

$$\mathbb{E}_t[r(I_t, J_t)] = r(\mathbf{p}'_t, J_t).$$

We note that by definition of the forecaster, $\mathbf{p}'_t = \mathbf{p}_n$ if t belongs to the n -th block. By a version of the Hoeffding-Azuma inequality for sums of Hilbert space-valued martingale differences stated as³ Lemma 3.2 in Chen and White (1996), we therefore get that with probability at least $1 - \delta$,

$$\left\| \frac{1}{NL} \sum_{t=1}^{NL} r(I_t, J_t) - \frac{1}{N} \sum_{n=1}^N r(\mathbf{p}_n, \hat{\mathbf{q}}_n) \right\|_2 \leq 4R \sqrt{\frac{\ln(2/\delta)}{NL}} \leq 4R \sqrt{\frac{\ln(2/\delta)}{T-L}}.$$

The second inequality comes from $NL = T - k \geq T - L$.

Step 3: Assertion (11). Since by definition $\mathbf{p}_n = (1 - \gamma) \mathbf{x}_n + \gamma \mathbf{u}$, we get

$$\left\| \frac{1}{N} \sum_{n=1}^N r(\mathbf{p}_n, \hat{\mathbf{q}}_n) - \frac{1}{N} \sum_{n=1}^N r(\mathbf{x}_n, \hat{\mathbf{q}}_n) \right\|_2 \leq 2\gamma R.$$

Step 4: Assertion (12). We fix a given block n . Lemma 23 indicates that with probability $1 - \delta$,

$$\left\| \hat{\sigma}_n - \tilde{H}(\hat{\mathbf{q}}_n) \right\|_2 \leq \sqrt{N_{\mathcal{I}} N_{\mathcal{H}}} \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta} \right). \quad (14)$$

Since Φ is Lipschitz (see Remark 20), with a Lipschitz constant in ℓ^2 -norms denoted by κ_{Φ} , we get that with probability $1 - \delta$,

$$\left\| \Phi(\hat{\sigma}_n) - \Phi(\tilde{H}(\hat{\mathbf{q}}_n)) \right\|_2 \leq \kappa_{\Phi} \sqrt{N_{\mathcal{I}} N_{\mathcal{H}}} \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta} \right).$$

3. Note that the martingale increments are bounded in norm by $2R$ in our framework and that $\sqrt{u} e^{-u} \leq e^{-u/2}$ for all $u \geq 0$.

By a union bound, the above bound holds for all blocks $n = 1, \dots, N$ with probability at least $1 - N\delta$. Finally, an application of Lemma 6 shows that

$$\frac{1}{N} \sum_{n=1}^N \bar{m} \left(\boldsymbol{\theta}_n, \Phi \left(\tilde{H}(\hat{\mathbf{q}}_n) \right) \right)$$

is in a ε_T -neighborhood (in ℓ^2 -norm) of

$$\frac{1}{N} \sum_{n=1}^N \bar{m} \left(\boldsymbol{\theta}_n, \Phi(\hat{\sigma}_n) \right),$$

where

$$\varepsilon_T = R\sqrt{N_{\mathcal{B}}} \left(\kappa_{\Phi} \sqrt{N_{\mathcal{I}}N_{\mathcal{H}}} \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}}N_{\mathcal{H}}}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2N_{\mathcal{I}}N_{\mathcal{H}}}{\delta} \right) \right).$$

Step 5: Assertion (13). Since \mathcal{C} is \bar{m} -approachable and by definition of the choices of the $\boldsymbol{\theta}_n$ in Figure 1, we get by Theorem 8, with probability 1,

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{N} \sum_{n=1}^N \bar{m} \left(\boldsymbol{\theta}_n, \Phi(\hat{\sigma}_n) \right) \right\|_2 \leq \frac{2R}{\sqrt{N}} \sqrt{N_{\mathcal{A}}N_{\mathcal{B}}} \leq \frac{2R}{\sqrt{T/L - 1}} \sqrt{N_{\mathcal{A}}N_{\mathcal{B}}},$$

since, as used already at the end of step 2, $N \geq T/L - 1$.

Conclusion of the proof. The proof is concluded by putting the pieces together, thanks to a triangle inequality. By a union bound, the obtained bound holds however only with probability at least $1 - (N + 1)\delta \geq 1 - ((T + L)/L)\delta$, where we used $N \leq T/L$. The stated bound follows by replacing all occurrences of δ in the previous steps by $\delta L/(T + L)$. ■

5.3.4 UNIFORM GUARANTEES OVER TIME

We present here a variant of the strategy of Figure 1 that r -approaches \mathcal{C} for the signaling structure H . This is achieved by making the strategy independent of the horizon T . (The strategy of the previous section depended on the knowledge of T , via suitable choices for L and γ .) Two options could have been worked out: resorting to some “doubling trick” or having the parameters L and γ vary over time. In the latter option, the lengths L_n of blocks n and the exploration rates γ_n used therein are no longer constant but of lengths polynomial in n . We chose the latter option for the sake of elegance and because it relies on a result of independent interest, namely a generalization of Theorem 3 to polynomial averages. We only state this generalization for mixed actions taken and observed, but the adaptation for pure actions follows easily.

Consider the setting of Theorem 3. The studied strategy relies on a parameter $\alpha \geq 0$. It plays an arbitrary \mathbf{x}_1 . For $t \geq 1$, it forms at stage $t + 1$ the vector-valued polynomial average

$$\hat{m}_t^\alpha = \frac{1}{T_t^\alpha} \sum_{s=1}^t s^\alpha m(\mathbf{x}_s, \mathbf{y}_s) \quad \text{where} \quad T_t^\alpha = \sum_{s=1}^t s^\alpha,$$

computes its projection c_t^α onto \mathcal{C} , and resorts to a mixed action \mathbf{x}_{t+1} solving the minimax equation

$$\min_{\mathbf{x} \in \Delta(\mathcal{A})} \max_{\mathbf{y} \in \Delta(\mathcal{B})} \langle \widehat{m}_t^\alpha - c_t^\alpha, m(\mathbf{x}, \mathbf{y}) - c_t^\alpha \rangle.$$

Theorem 25 *We denote by M a bound in norm over m , i.e.,*

$$\max_{(a,b) \in \mathcal{A} \times \mathcal{B}} \|m(a, b)\|_2 \leq M.$$

For all $\alpha \geq 0$, when \mathcal{C} is an approachable closed convex set, the above strategy ensures that for all strategies of the second player, with probability 1, for all $T \geq 1$,

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{\sum_{t=1}^T t^\alpha} \sum_{t=1}^T t^\alpha m(\mathbf{x}_t, \mathbf{y}_t) \right\|_2 \leq 2M \frac{\sqrt{\sum_{t=1}^T t^{2\alpha}}}{\sum_{t=1}^T t^\alpha} \leq \frac{2M(\alpha + 1)}{\sqrt{T}}. \quad (15)$$

It is interesting to note that the convergence rate is independent of α and is the same as standard approachability ($1/\sqrt{T}$). The proof of this theorem is a slight modification of the proof of Theorem 3 and is hence deferred to Appendix D.

The extension to polynomially weighted averages can also be obtained in the context of set-valued approachability. This is because the key to Theorem 8 is Lemma 10, which indicates that to get set-valued approachability, it suffices to approach, in the usual sense, $\widetilde{\mathcal{C}}$. Both can thus be performed with polynomially weighted averages.

Consider now the variant of the strategy of Figure 1 for which the length of the n -th block, denoted by L_n , is equal to n^α , the exploration rate on this block comes at a rate $\gamma_n = n^{-\alpha/3}$, and Ψ is an \overline{m} -approaching strategy of \mathcal{C} with respect to polynomially weighted averages with parameter $\alpha = 3/2$. We call it a time-adaptive version of the strategy of Figure 1; indeed, this choice of α ensures that there are roughly $T^{2/5}$ blocks and that the length of the last one is of the order of $T^{3/5}$. Note that it does not depend anymore on any time horizon T , hence guarantees can be obtained for all T .

Theorem 26 *In the same setting and under the same assumptions as in Theorem 24, the time-adaptive version of the strategy described in Figure 1 (with $L_n = n^\alpha$ and $\gamma_n = n^{-\alpha/3}$ for $\alpha = 3/2$) ensures that, for all strategies of the second player, for all $T \geq 1$ and all $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2 \leq \Xi \left(T^{-1/5} \sqrt{\ln \frac{T}{\delta}} + T^{-2/5} \ln \frac{T}{\delta} \right)$$

for some constant Ξ depending only on \mathcal{C} and on the game (r, H) at hand.

The proof follows closely the one of Theorem 24 and is presented in Appendix E.

Corollary 27 *In the same setting and under the same assumptions as in Theorem 24, the time-adaptive version of the strategy described in Figure 1 (with $L_n = n^\alpha$ and $\gamma_n = n^{-\alpha/3}$ for $\alpha = 3/2$) indeed r -approaches \mathcal{C} for the signalling structure H .*

Proof The strategy at hand is such that for all $T \geq 1$, with probability at least $1 - 1/T^2$,

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2 \leq \Xi(T^{-1/5} \ln T^3 + T^{-2/5} \ln T^3).$$

In particular, a union bound shows that for all $T \geq 2$,

$$\sup_{\tau \geq T} \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{\tau} \sum_{t=1}^{\tau} r(I_t, J_t) \right\|_2 \leq R_T \stackrel{\text{def}}{=} \sup_{\tau \geq T} \Xi(\tau^{-1/5} \ln \tau^3 + \tau^{-2/5} \ln \tau^3),$$

with probability at least $1 - P_T$, where $P_T = \sum_{\tau \geq T} 1/\tau^2$. We note that $P_T \rightarrow 0$ and $R_T \rightarrow 0$ as $T \rightarrow \infty$. To see that the definition of approachability is satisfied, given $\varepsilon > 0$, it suffices to define T_ε as the minimal T such that $R_T \leq \varepsilon$ and $P_T \leq \varepsilon$. ■

5.4 Approachability in the Case of General Games with Partial Monitoring

Unfortunately, as is illustrated in the example below, there exist games with partial monitoring that are not bi-piecewise linear. However, we will show that if Condition (APM) holds there exist strategies with a constant per-round complexity that approach polytopes even when the game is not bi-piecewise linear. That is, by considering simpler closed convex sets \mathcal{C} , no assumption is needed on the pair (r, H) .

We will conclude this main part of the paper by re-proving, using a doubling trick, that Condition (APM) is still sufficient for approachability in the most general case when no assumption is made neither on (r, H) nor on \mathcal{C} , at the cost of inefficiency.

Example 3 *The following game (with the same action and signal sets as in Example 2) is not bi-piecewise linear.*

	L	M	R
T	$(1, 0, 0, 0) / \clubsuit$	$(0, 0, 1, 0) / \clubsuit$	$(2, 0, 4, 0) / \heartsuit$
B	$(0, 1, 0, 0) / \clubsuit$	$(0, 0, 0, 1) / \clubsuit$	$(0, 3, 0, 5) / \heartsuit$

Proof We denote mixed actions of the first player by $(p, 1 - p)$, where $p \in [0, 1]$ denotes the probability of playing T and $1 - p$ is the probability of playing B . It is immediate that $\bar{m}((p, 1 - p), \clubsuit)$ can be identified with the set of all product distributions on 2×2 elements with first marginal distribution $(p, 1 - p)$. The proof of Lemma 19 shows that the set \mathcal{B} associated with any game always contains the Dirac masses on each signal; that is, $\delta_{\clubsuit} \in \mathcal{B}$. But for $p \neq p'$ and $\lambda \in (0, 1)$, denoting $\bar{p} = \lambda p + (1 - \lambda)p'$, one necessarily has that

$$\bar{m}((\bar{p}, 1 - \bar{p}), \clubsuit) \not\subseteq \lambda \bar{m}((p, 1 - p), \clubsuit) + (1 - \lambda) \bar{m}((p', 1 - p'), \clubsuit);$$

the inclusion \subseteq holds by concavity of \bar{m} in its first argument (Lemma 17) but this inclusion is always strict here since the left-hand side is formed by product distributions while the right-hand side also contains distributions with correlations. Hence, bi-piecewise linearity cannot hold for this game. ■

5.4.1 APPROACHABILITY OF THE NEGATIVE ORTHANT IN GENERAL GAMES

For the sake of simplicity, we start with the case of the negative orthant \mathbb{R}_-^d and prove the following result. Note that for the first time in this general framework, the rates for approachability of polytopes are independent of the dimension (as is the case with full monitoring).

Theorem 28 *If Condition (APM) is satisfied for \bar{m} and \mathbb{R}_-^d , then there exists a strategy for (r, H) -approaching \mathbb{R}_-^d at a rate of the order of $T^{-1/5}$, with a constant per-round complexity.*

Our argument is based on Lemma 19; we use in the sequel the objects and notation introduced therein. We denote by $r = (r_k)_{1 \leq k \leq d}$ the components of the d -dimensional payoff function r and introduce, for all $k \in \{1, \dots, d\}$, the set-valued mapping \tilde{m}_k defined by

$$\tilde{m}_k : (\mathbf{p}, b) \in \Delta(\mathcal{I}) \times \mathcal{B} \mapsto \tilde{m}_k(\mathbf{p}, b) = \left\{ r_k(\mathbf{p}, \mathbf{q}) : \mathbf{q} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}) = b \right\}.$$

The mapping \tilde{m} is then defined as the Cartesian product of the \tilde{m}_k ; formally, for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $b \in \mathcal{B}$,

$$\tilde{m}(\mathbf{p}, b) = \left\{ (z_1, \dots, z_d) : \forall k \in \{1, \dots, d\}, z_k \in \tilde{m}_k(\mathbf{p}, b) \right\}.$$

We then linearly extend this mapping to a set-valued mapping \tilde{m} defined on $\Delta(\mathcal{I}) \times \Delta(\mathcal{B})$ and finally consider the set-valued mapping \check{m} defined on $\Delta(\mathcal{I}) \times \mathcal{F}$ by

$$\forall \sigma \in \mathcal{F}, \quad \forall \mathbf{p} \in \Delta(\mathcal{I}), \quad \check{m}(\mathbf{p}, \sigma) = \tilde{m}(\mathbf{p}, \Phi(\sigma)) = \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \tilde{m}(\mathbf{p}, b),$$

where Φ refers to the mapping defined in Lemma 19 (based on \bar{m}). The lemma below indicates why \check{m} is an excellent substitute to \bar{m} in the case of the approachability of the orthant \mathbb{R}_-^d .

Lemma 29 *The set-valued mappings \check{m} and \bar{m} satisfy that for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$,*

1. *the inclusion $\bar{m}(\mathbf{p}, \sigma) \subseteq \check{m}(\mathbf{p}, \sigma)$ holds;*
2. *if $\bar{m}(\mathbf{p}, \sigma) \subseteq \mathbb{R}_-^d$, then one also has $\check{m}(\mathbf{p}, \sigma) \subseteq \mathbb{R}_-^d$.*

The interpretation of these two properties is: 1. \check{m} -approaching a set \mathcal{C} is more difficult than \bar{m} -approaching it; and 2. that if Condition (APM) holds for \bar{m} and \mathbb{R}_-^d , it also holds for \check{m} and \mathbb{R}_-^d .

Proof For the first property, note that by the component-wise construction of \tilde{m} ,

$$\forall b \in \mathcal{B}, \quad \forall \mathbf{p} \in \Delta(\mathcal{I}), \quad \bar{m}(\mathbf{p}, b) \subseteq \tilde{m}(\mathbf{p}, b).$$

Lemma 19, the linear extension of \tilde{m} , and the definition of \check{m} then show that

$$\forall \sigma \in \mathcal{F}, \quad \forall \mathbf{p} \in \Delta(\mathcal{I}), \quad \bar{m}(\mathbf{p}, \sigma) = \sum_{b \in \mathcal{B}} \Phi_b(\sigma) \bar{m}(\mathbf{p}, b) \subseteq \tilde{m}(\mathbf{p}, \Phi(\sigma)) = \check{m}(\mathbf{p}, \sigma).$$

As for the second property, it suffices to work component-wise. Note that (by Lemma 19 again) the stated assumption exactly means that $\sum_{b \in \mathcal{B}} \Phi_b(\sigma) \bar{m}(\mathbf{p}, b) \subset \mathbb{R}_-^d$. In particular, rewriting the non-positivity constraint for each of the d components of the payoff vectors, we get

$$\sum_{b \in \mathcal{B}} \Phi_b(\sigma) \tilde{m}_k(\mathbf{p}, b) \subseteq \mathbb{R}_-,$$

for all $k \in \{1, \dots, d\}$; thus, in particular, $\sum_{b \in \mathcal{B}} \Phi_b(\sigma) \tilde{m}(\mathbf{p}, b) = \check{m}(\mathbf{p}, \sigma) \subseteq \mathbb{R}_-^d$. ■

We can then extend the result of the previous section without the bi-piecewise linearity assumption.

Proof [of Theorem 28] The assumption of the theorem and the second property of Lemma 29 imply that Condition (APM) holds for \mathbb{R}_-^d and \check{m} . Furthermore, the latter corresponds to a bi-piecewise linear game, i.e., Assumption 1 is satisfied. Indeed, we show below that each $\tilde{m}_k(\cdot, b)$ is a piecewise linear function. As a consequence, each $\check{m}(\cdot, b)$ is also a piecewise linear function.

Each $\tilde{m}_k(\cdot, b)$ is piecewise linear since \tilde{m}_k is based on the scalar payoff function r_k . Indeed, since \tilde{H} is linear, the set

$$\left\{ \mathbf{q} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}) = b \right\}$$

is a polytope, thus, the convex hull of some finite set $\{\mathbf{q}_{b,1}, \dots, \mathbf{q}_{b,M}\} \subset \Delta(\mathcal{J})$. Therefore, for every $\mathbf{p} \in \Delta(\mathcal{I})$, by linearity of r_k (and by the fact that it takes one-dimensional values),

$$\begin{aligned} \tilde{m}_k(\mathbf{p}, b) &= \left\{ r_k(\mathbf{p}, \mathbf{q}) : \mathbf{q} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}) = b \right\} \\ &= \text{co} \left\{ r_k(\mathbf{p}, \mathbf{q}_{b,1}), \dots, r_k(\mathbf{p}, \mathbf{q}_{b,M}) \right\} = \left[\min_{k \in \{1, \dots, M\}} r_k(\mathbf{p}, \mathbf{q}_{b,k}), \max_{k' \in \{1, \dots, M\}} r_k(\mathbf{p}, \mathbf{q}_{b,k'}) \right], \end{aligned} \quad (16)$$

where co stands for the convex hull. Since all mappings $r_k(\cdot, \mathbf{q}_{b,k})$ are linear, their minimum and their maximum are piecewise linear functions, therefore $\tilde{m}_k(\cdot, b)$ is also piecewise linear.

Therefore, the steps between Equations (11)–(13) of the proof of Theorem 24 (or the corresponding statements in the proof of Theorem 26) can be adapted by replacing \bar{m} and \check{m} by, respectively, \tilde{m} , \check{m} , and its extension corresponding to Definition 21. The result about approachability rates follows.

We now prove that the strategy constructed here is efficient. Indeed, recall that this is always the case as long as the projection onto the associated convex set $\tilde{\mathcal{C}}$ defined by (4) with the linear function \bar{m} is also efficient. But this follows from the fact that as \mathbb{R}_-^d is a polyhedron, the set $\tilde{\mathcal{C}}$ is a polytope. ■

We now generalize the above ideas to more complex sets.

5.4.2 APPROACHABILITY OF POLYTOPES FOR GENERAL GAMES

If the target set \mathcal{C} is a polytope, then \mathcal{C} can be written as the intersection of a finite number of half-planes, i.e., there exists a finite family $\{(e_k, f_k) \in \mathbb{R}^d \times \mathbb{R}, k \in \mathcal{K}\}$ such that

$$\mathcal{C} = \{z \in \mathbb{R}^d : \langle z, e_k \rangle \leq f_k, \forall k \in \mathcal{K}\}.$$

Given the original (not necessarily bi-piecewise linear) game (r, H) , we introduce another game $(r_{\mathcal{C}}, H)$, whose payoff function $r_{\mathcal{C}} : \mathcal{I} \times \mathcal{J} \rightarrow \mathbb{R}^{\mathcal{K}}$ is defined as

$$\forall i \in \mathcal{I}, \quad \forall j \in \mathcal{J}, \quad r_{\mathcal{C}}(i, j) = \left[\langle r(i, j), e_k \rangle - f_k \right]_{k \in \mathcal{K}}.$$

The following lemma is a mere exercise of rewriting.

Lemma 30 *Given a polytope \mathcal{C} , the (r, H) -approachability of \mathcal{C} and the $(r_{\mathcal{C}}, H)$ -approachability of \mathbb{R}^d are equivalent in the sense that every strategy for solving one problem translates to a strategy for solving the other problem. In addition, Condition (APM) holds for (r, H) and \mathcal{C} if and only if it holds for $(r_{\mathcal{C}}, H)$ and \mathbb{R}^d .*

Via the lemma above, Theorem 28 indicates that Condition (APM) for (r, H) and \mathcal{C} is a sufficient condition for the (r, H) -approachability of \mathcal{C} and provides an efficient strategy to do so. (The per-round complexity of this strategy depends in particular at least linearly on the cardinality of \mathcal{K} .) Again, rates of convergence are also, for the first time, independent of the dimensions (yet the question of their optimality remains open).

5.4.3 APPROACHABILITY OF GENERAL CONVEX SETS IN THE CASE OF GENERAL GAMES

In the above, we provided efficient strategies in the following cases:

- Up to projection oracles, when the games (r, H) are bi-piecewise linear, with no assumption on the target set \mathcal{C} ; see Section 5.3. This includes at least the minimization of external and internal regret, for which the projections can indeed be performed efficiently; see the upcoming Section 6.
- When the target set \mathcal{C} is a polytope, with no assumption on the game (r, H) ; see Sections 5.4.1 and 5.4.2.

We only mention the case of general games (r, H) and general closed convex target sets \mathcal{C} in this section to have a complete, self-contained, and constructive proof of the sufficiency of Condition (APM) for (r, H) -approachability. (Perchet, 2011a already proved the latter.)

Theorem 15 and Lemma 17 show that Condition (APM) is indeed sufficient to (r, H) -approach any general closed convex set \mathcal{C} . However, the computational complexity of the resulting strategy is much larger as the per-round complexity increases over time. Another way to deal with a general closed convex set is based on the fact that it can be approximated arbitrarily well by a polytope (where the number of vertices of the latter increases as the quality of the approximation does). Playing in regimes approachability strategies of such a sequence of approximations also gives an approachability strategy of the original set \mathcal{C} . However, the per-round complexity increases over time (as the numbers of vertices of the approximating polytopes do).

6. Application to Regret Minimization

In this section we analyze external and internal regret minimization in repeated games with partial monitoring from the perspective of approachability. We show how to efficiently minimize regret in both setups using the results developed for vector-valued games with partial monitoring. To do so, we indicate why the assumption of bi-piecewise linearity (Assumption 1) is satisfied.

The results instantiated below are not necessarily new in terms of efficiency or convergence rates, and some are even slightly suboptimal. However, our point is that all previous good strategies were specifically designed for the problem of regret minimization, while we introduced above a general strategy for all approachability problems, including regret minimization. And what we gained in generality (the wider range of problems that we can deal with) has no impact (or little impact only) on the efficiency or on the rates, which we think is an important contribution.

6.1 External Regret

We consider in this section the framework and aim introduced by Rustichini (1999) and studied, sometimes for restricted classes of games, by Piccolboni and Schindelhauer (2001), Mannor and Shimkin (2003), Cesa-Bianchi et al. (2006), Lugosi et al. (2008), Bartók et al. (2010, 2011), Foster and Rakhlin (2012). We show that our general strategy can be used for regret minimization.

Scalar payoffs are obtained (but not observed) by the first player, i.e., $d = 1$: the payoff function r is a mapping $\mathcal{I} \times \mathcal{J} \rightarrow \mathbb{R}$. We still denote by R a bound on $|r|$. We define in this section

$$\widehat{\mathbf{q}}_T = \frac{1}{T} \sum_{t=1}^T \delta_{J_t}$$

as the empirical distribution of the actions taken by the second player during the first T rounds. (This is in slight contrast with the notation $\widehat{\mathbf{q}}_n$ used in Section 5.3 to denote such an empirical distribution, but only taken within regime n .)

The external regret of the first player at round T equals by definition

$$R_T^{\text{ext}} = \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho\left(\mathbf{p}, \widetilde{H}(\widehat{\mathbf{q}}_T)\right) - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t),$$

where $\rho : \Delta(\mathcal{I}) \times \mathcal{F}$ is defined as follows: for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$,

$$\rho(\mathbf{p}, \sigma) = \min \left\{ r(\mathbf{p}, \mathbf{q}) : \mathbf{q} \text{ such that } \widetilde{H}(\mathbf{q}) = \sigma \right\}. \quad (17)$$

The function ρ is continuous in its first argument and therefore the supremum in the defining expression of R_T^{ext} is a maximum.

We recall briefly why, intuitively, this is the natural notion of external regret to consider in this case (more formal arguments are given in Rustichini, 1999). Indeed, the first term in the definition of R_T^{ext} is (close to) the worst-case average payoff obtained by the first player when playing consistently a mixed action \mathbf{p} against a sequence of mixed actions inducing on average the same laws on the signals as the sequence of actions actually played.

Rustichini (1999) calls the partial monitoring in the game (r, H) statistically sufficient when

$$\max_{\mathbf{p} \in \Delta(\mathcal{I})} r(\mathbf{p}, \hat{\mathbf{q}}_T) = \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \tilde{H}(\hat{\mathbf{q}}_T)).$$

In general, only an inequality \geq holds between the two quantities. A line of research initiated by Piccolboni and Schindelhauer (2001) first studied efficient strategies to minimize the regret in the said case of a statistically sufficient monitoring.

6.1.1 A STRATEGY MINIMIZING EXTERNAL REGRET

The following result is a consequence of Theorem 26, as its proof shows; it corresponds to the main result of Lugosi et al. (2008), with the same convergence rate but with a different strategy. (However, Perchet, 2011b, Section 2.3 exhibited an efficient strategy achieving a convergence rate of order $T^{-1/3}$, which is optimal; this strategy was an ad hoc strategy for regret minimization. Nonetheless, a question that remains open is thus whether the rates exhibited in Theorem 26 could be improved.)

Corollary 31 *The first player can apply the strategy of Theorem 25 such that for all strategies of the second player, for all T and all $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$R_T^{\text{ext}} \leq \Xi \left(T^{-1/5} \sqrt{\ln \frac{T}{\delta}} + T^{-2/5} \ln \frac{T}{\delta} \right)$$

for some constant Ξ depending only on the game (r, H) at hand.

The discussion about the efficiency of the strategy is postponed to the end of this section, as it relies on some objects that will be introduced in the proof of the corollary. The latter proof is an extension to the setting of partial monitoring of the original proof and strategy of Blackwell (1956b) for the case of external regret under full monitoring: in the latter case the vector-payoff function \underline{r} and the set \mathcal{C} considered in our proof are equal to the ones considered by Blackwell.

Proof We embed \mathcal{F} into $\mathbb{R}^{\mathcal{I} \times \mathcal{H}}$ so that in this proof we will be working in the vector space $\mathbb{R}^d = \mathbb{R} \times \mathbb{R}^{\mathcal{I} \times \mathcal{H}}$. We consider the closed convex set \mathcal{C} and the vector-valued payoff function \underline{r} respectively defined by

$$\mathcal{C} = \left\{ (z, \sigma) \in \mathbb{R} \times \mathcal{F} : z \geq \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \sigma) \right\} \quad \text{and} \quad \underline{r}(i, j) = \begin{bmatrix} r(i, j) \\ \tilde{H}(\delta_j) \end{bmatrix},$$

for all $(i, j) \in \mathcal{I} \times \mathcal{J}$.

We first show that Condition (APM) is satisfied for the considered convex set \mathcal{C} and game (\underline{r}, H) . To do so, by continuity of ρ in its first argument, we associate with each $\mathbf{q} \in \Delta(\mathcal{J})$ an element $\phi(\mathbf{q}) \in \Delta(\mathcal{I})$ such that

$$\phi(\mathbf{q}) \in \operatorname{argmax}_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \tilde{H}(\mathbf{q})).$$

Then, given any $\mathbf{q} \in \Delta(\mathcal{J})$, we note that for all \mathbf{q}' satisfying $\tilde{H}(\mathbf{q}') = \tilde{H}(\mathbf{q})$, we have by definition of ρ ,

$$r(\phi(\mathbf{q}), \mathbf{q}') \geq \rho(\phi(\mathbf{q}), \tilde{H}(\mathbf{q}')) = \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \tilde{H}(\mathbf{q}')),$$

which shows that $\underline{r}(\phi(\mathbf{q}), \mathbf{q}') \in \mathcal{C}$. The required condition is thus satisfied.

We then show that Assumption 1 is satisfied. To do so, we will use the same arguments as around (16) and actually prove the stronger property that the mappings $\overline{m}(\cdot, \sigma)$ are piecewise linear for all $\sigma \in \mathcal{F}$; we fix such a σ in the sequel. Only the first coordinate r of \underline{r} depends on \mathbf{p} , so the desired property is true if and only if the mapping $\overline{m}_1(\cdot, \sigma)$ defined by

$$\mathbf{p} \in \Delta(\mathcal{I}) \longmapsto \overline{m}_1(\mathbf{p}, \sigma) = \left\{ r(\mathbf{p}, \mathbf{q}) : \mathbf{q} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}) = \sigma \right\}$$

is piecewise linear. But this is true because r takes scalar values, as indicated around (16).

Theorem 26 can then be applied to exhibit the convergence rates; we simply need to relate the quantity of interest here to the one considered therein. To that end we use the fact that the mapping

$$\sigma \in \mathcal{F} \longmapsto \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \sigma)$$

is Lipschitz, with Lipschitz constant in ℓ^2 -norm denoted by L_ρ ; the proof of this fact is detailed in the last paragraph of this proof.

Now, the regret is non-positive if $\sum_{t=1}^T \underline{r}(I_t, J_t)/T$ belongs to \mathcal{C} ; we therefore only need to consider the case when this average is not in \mathcal{C} . In the latter case, we denote by $(\tilde{r}_T, \tilde{\sigma}_T)$ its projection in ℓ^2 -norm onto \mathcal{C} . We have first that the defining inequality of \mathcal{C} is an equality on its border, so that

$$\tilde{r}_T = \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \tilde{\sigma}_T);$$

and second, that

$$\begin{aligned} R_T^{\text{ext}} &= \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \tilde{H}(\hat{\mathbf{q}}_T)) - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \\ &\leq \left| \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \tilde{H}(\hat{\mathbf{q}}_T)) - \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \tilde{\sigma}_T) \right| + \left| \tilde{r}_T - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right| \\ &\leq L_\rho \left\| \tilde{\sigma}_T - \tilde{H}(\hat{\mathbf{q}}_T) \right\|_2 + \left| \tilde{r}_T - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right| \\ &\leq \sqrt{2} \max\{L_\rho, 1\} \left\| \begin{bmatrix} \tilde{r}_T \\ \tilde{\sigma}_T \end{bmatrix} - \frac{1}{T} \sum_{t=1}^T \underline{r}(I_t, J_t) \right\|_2 \\ &= \sqrt{2} \max\{L_\rho, 1\} \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T \underline{r}(I_t, J_t) \right\|_2. \end{aligned}$$

The claimed rates are now seen to follow from the ones indicated in Theorem 26.

It only remains to prove the indicated Lipschitz continuity. (All Lipschitz continuity statements that follow will be with respect to the ℓ^2 -norms.) We have by Definition 21 that for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$,

$$\rho(\mathbf{p}, \sigma) = \min \overline{m}_1(\mathbf{p}, \Phi(\sigma)),$$

where the linear \overline{m}_1 indifferently either is relative to \overline{m}_1 or is the projection onto the first component of the function \overline{m} relative to \overline{m} . By Remark 20 the mapping $\sigma \in \mathcal{F} \mapsto \Phi(\sigma)$ is κ_Φ -Lipschitz; this entails, by Lemma 6, that for all $\mathbf{p} \in \Delta(\mathcal{I})$, the mapping $\sigma \in \mathcal{F} \mapsto \rho(\mathbf{p}, \sigma)$ is $R\sqrt{N_{\mathcal{B}}}$ κ_Φ -Lipschitz. In particular, since the latter Lipschitz constant is independent of \mathbf{p} , the mapping

$$\sigma \in \mathcal{F} \mapsto \max_{\mathbf{p} \in \Delta(\mathcal{I})} \rho(\mathbf{p}, \sigma)$$

is $R\sqrt{N_{\mathcal{B}}}$ κ_Φ -Lipschitz as well, which concludes the proof. ■

6.1.2 DISCUSSION ABOUT EFFICIENCY

An argument similar to the one in Perchet (2011b) shows that the convex set \mathcal{C} is defined by a finite number of piecewise linear equations, it is therefore a polyhedron; so that the projection onto it, as well as the computation of the strategy, can be done efficiently. We only sketch here the argument. The argument used when referring to (16) indicates a priori that for each $\sigma \in \mathcal{F}$, there exist a finite number K_σ (depending on σ) of mixed actions $\mathbf{q}_{\sigma,1}, \dots, \mathbf{q}_{\sigma,M_\sigma}$ such that for all $\mathbf{p} \in \Delta(\mathcal{I})$, we have $\rho(\mathbf{p}, \sigma) = \min\{r(\mathbf{p}, \mathbf{q}_{\sigma,1}), \dots, r(\mathbf{p}, \mathbf{q}_{\sigma,M_\sigma})\}$. But by an argument stated in Perchet (2011b),

$$\sigma \mapsto \left\{ \mathbf{q} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}) = \sigma \right\}$$

evolves in a piecewise linear way and thus there exist a finite number K of piecewise linear functions $\sigma \mapsto q'_{\sigma,k}$, with $k = 1, \dots, K$, such that, for all $\sigma \in \mathcal{F}$,

$$\{\mathbf{q}_{\sigma,1}, \dots, \mathbf{q}_{\sigma,K_\sigma}\} = \{q'_{\sigma,1}, \dots, q'_{\sigma,K}\}.$$

(There can be some redundancies between the $q'_{\sigma,k}$.) Because of this, we have that for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$,

$$\rho(\mathbf{p}, \sigma) = \min\{r(\mathbf{p}, q'_{\sigma,1}), \dots, r(\mathbf{p}, q'_{\sigma,K})\}.$$

Each function $\sigma \mapsto q'_{\sigma,k}$ being piecewise linear, one can construct a finite set $\{\mathbf{p}_1, \dots, \mathbf{p}_{\tilde{K}}\} \subset \Delta(\mathcal{I})$ such that, for any $\sigma \in \mathcal{F}$, the mapping $\mathbf{p} \mapsto \rho(\mathbf{p}, \sigma)$ is maximized at one of these \mathbf{p}_k . The convex set \mathcal{C} is therefore defined by a finite number of piecewise linear equations, it is a polyhedron. Its lifted image $\tilde{\mathcal{C}}$ is then a polytope: thus, the projection onto it, hence the computation of the proposed strategy, can be done efficiently.

6.2 Internal / Swap Regret

Foster and Vohra (1999) defined internal regret with full monitoring as follows. A player has no internal regret if, for every action $i \in \mathcal{I}$, he has no external regret on the stages when this specific action i was played (if there are enough such stages). In other words, i is the best response to the empirical distribution of actions of the other player on these stages.

With partial monitoring, the first player evaluates his payoffs in a pessimistic way through the function ρ defined in (17). This function is not linear over $\Delta(\mathcal{I})$ in general (it is concave), so that the best responses are not necessarily pure actions $i \in \mathcal{I}$ but mixed

actions, i.e., elements of $\Delta(\mathcal{I})$. Following Lehrer and Solan (2007) one therefore can partition the stages not depending on the pure actions actually played but on the mixed actions $\mathbf{p}_t \in \Delta(\mathcal{I})$ used to draw them. To this end, it is convenient to assume that the strategies of the first player need to pick these mixed actions in a finite grid of $\Delta(\mathcal{I})$, which we denote by $\{\mathbf{p}_g, g \in \mathcal{G}\}$, where \mathcal{G} is a finite set. At each round t , the first player picks an index $G_t \in \mathcal{G}$ and uses the distribution \mathbf{p}_{G_t} to draw his action I_t . A discussion about the choice of \mathcal{G} is provided below. For now, we define formally \mathcal{G} -internal regret as internal regret with respect to the set of mixed actions \mathcal{G} .

Up to a standard concentration-of-the-measure argument, we will measure the payoff at round t with $r(\mathbf{p}_{G_t}, J_t)$ rather than with $r(I_t, J_t)$. For each $g \in \mathcal{G}$, we denote by $N_T(g)$ the number of stages in $\{1, \dots, T\}$ for which we had $G_t = g$ and, whenever $N_T(g) > 0$,

$$\hat{\mathbf{q}}_{T,g} = \frac{1}{N_T(g)} \sum_{t:G_t=g} \delta_{J_t}.$$

We define $\hat{\mathbf{q}}_{T,g}$ in an arbitrary way when $N_T(g) = 0$. The \mathcal{G} -internal regret of the first player at round T is measured as

$$R_T^{\text{int}} = \max_{g,g' \in \mathcal{G}} \frac{N_T(g)}{T} \left(\rho(\mathbf{p}_{g'}, \tilde{H}(\hat{\mathbf{q}}_{T,g})) - r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \right).$$

Actually, our proof technique rather leads to the minimization of some \mathcal{G} -swap regret (see Blum and Mansour, 2007, for the definition of swap regret in full monitoring):

$$R_T^{\text{swap}} = \sum_{g \in \mathcal{G}} \frac{N_T(g)}{T} \left(\max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \tilde{H}(\hat{\mathbf{q}}_{T,g})) - r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \right)_+.$$

At first sight, to handle all possible alternatives one should take \mathcal{G} as a thin grid in $\Delta(\mathcal{I})$, i.e., some ε -discretization of the latter. This is what Lehrer and Solan (2007) do. However, Perchet (2011b) showed that there exists a finite subset \mathcal{G}_0 of $\Delta(\mathcal{I})$ such that \mathcal{G}_0 contains a best response to any mixed action of the second player: for all $\mathbf{q} \in \Delta(\mathcal{J})$,

$$\left(\underset{\mathbf{p} \in \Delta(\mathcal{I})}{\text{argmin}} \rho(\mathbf{p}, \tilde{H}(\mathbf{q})) \right) \cap \mathcal{G}_0 \neq \emptyset.$$

The strategy we discuss below will have a complexity polynomial in the size of \mathcal{G} . We thus advise to take $\mathcal{G} = \mathcal{G}_0$ for the sake of efficiency.

Again, the following bound on the swap regret easily follows from Theorem 24. The latter constructs a simple and direct strategy to control the swap regret, thus also the internal regret. It therefore improves on the results of Lehrer and Solan (2007) and Perchet (2009, 2011b), three papers that presented more involved and less efficient strategies to do so. These strategies were indeed based on auxiliary strategies using thin grids that need to be refined over time; this resulted in complexities that were at least exponential in the number of rounds. (The ideas used therein bear some resemblance with what is done in calibration, see the references provided in Section 4.) In contrast, our strategy can have a constant per-round complexity (when used with the grid \mathcal{G}_0). This is a major improvement in efficiency. However, as far as convergence rates are concerned, we must note that again, as in the case of external regret, Perchet (2011b) obtained rates of the faster order $T^{-1/3}$, for an ad hoc (inefficient) strategy. We thus sacrifice efficiency for rates.

Corollary 32 *The first player has an explicit strategy such that for all strategies of the second player, for all T and all $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$R_T^{\text{swap}} \leq \Xi \left(T^{-1/5} \sqrt{\ln \frac{T}{\delta}} + T^{-2/5} \ln \frac{T}{\delta} \right)$$

for some constant Ξ depending only on the game (r, H) at hand and on the size of the finite grid \mathcal{G} .

Proof The proof of this corollary is based on ideas similar to the ones used in the proof of Corollary 31; \mathcal{G} will play the role of the action set of the first player. The proof proceeds in four steps. In the first step, we construct an approachability setup and show that Condition (APM) applies. In the second step, we show that Assumption 1 is satisfied. In the third step we analyze the convergence rates of the swap regret. In the fourth and final step, we show that the set we are approaching possesses some smoothness properties by providing a uniform Lipschitz bound on certain functions.

Step 1: We denote by

$$\mathcal{F}_{\text{cone}} = \{ \lambda \sigma, \sigma \in \mathcal{F}, \lambda \in \mathbb{R}_+ \}$$

the cone generated by \mathcal{F} and extend linearly $\rho : \Delta(\mathcal{I}) \times \mathcal{F} \rightarrow \mathbb{R}$ into a mapping $\rho : \Delta(\mathcal{I}) \times \mathcal{F}_{\text{cone}} \rightarrow \mathbb{R}$ as follows: for all $\mathbf{p} \in \Delta(\mathcal{I})$, for all $\lambda \geq 0$ with $\lambda \neq 1$, and all $\sigma \in \mathcal{F}$,

$$\rho(\mathbf{p}, \lambda \sigma) = \begin{cases} 0 & \text{if } \lambda = 0, \\ \lambda \rho(\mathbf{p}, \sigma) & \text{if } \lambda > 0. \end{cases}$$

In the sequel, we embed $\mathcal{F}_{\text{cone}}$ into $\mathbb{R}^{\mathcal{I} \times \mathcal{H}}$.

The closed convex set \mathcal{C} and the vector-valued payoff function \underline{r} are then respectively defined by

$$\mathcal{C} = \left\{ (z_g, \mathbf{v}_g)_{g \in \mathcal{G}} \in (\mathbb{R} \times \mathcal{F}_{\text{cone}})^{\mathcal{G}} : \forall g \in \mathcal{G}, z_g \geq \max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \mathbf{v}_g) \right\}$$

and, for all $(g, j) \in \mathcal{G} \times \mathcal{J}$,

$$\underline{r}(g, j) = \left[\begin{array}{c} r(\mathbf{p}_g, j) \mathbb{I}_{\{g'=g\}} \\ \tilde{H}(\delta_j) \mathbb{I}_{\{g'=g\}} \end{array} \right]_{g' \in \mathcal{G}}.$$

To show that \mathcal{C} is \underline{r} -approachable, we associate with each $\mathbf{q} \in \Delta(\mathcal{J})$ an element $g^*(\mathbf{q}) \in \mathcal{G}$ such that

$$g^*(\mathbf{q}) \in \operatorname{argmax}_{g \in \mathcal{G}} \rho(\mathbf{p}_g, \tilde{H}(\mathbf{q})).$$

Then, given any $\mathbf{q} \in \Delta(\mathcal{J})$, we note that for all \mathbf{q}' satisfying $\tilde{H}(\mathbf{q}') = \tilde{H}(\mathbf{q})$, the components of the vector $\underline{r}(g^*(\mathbf{q}), \mathbf{q}')$ are all null but the ones corresponding to $g^*(\mathbf{q})$, for which we have

$$\begin{aligned} & r(\mathbf{p}_{g^*(\mathbf{q})}, \mathbf{q}') \\ & \geq \rho(\mathbf{p}_{g^*(\mathbf{q})}, \tilde{H}(\mathbf{q}')) = \rho(\mathbf{p}_{g^*(\mathbf{q})}, \tilde{H}(\mathbf{q})) = \max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \tilde{H}(\mathbf{q})) = \max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \tilde{H}(\mathbf{q}')), \end{aligned}$$

where the first inequality is by definition of ρ . Therefore, $\underline{r}(g^*(\mathbf{q}), \mathbf{q}') \in \mathcal{C}$. Condition (APM) in Lemma 22 and Theorem 24 is thus satisfied, so that we have approachability.

Step 2: We then show that Assumption 1 is satisfied. It suffices to show that for all $\sigma \in \mathcal{F}$, the mapping

$$\pi = (\pi_g)_{g \in \mathcal{G}} \in \Delta(\mathcal{G}) \longmapsto \bar{m}_1(\pi, \sigma) = \left\{ (\pi_g r(\mathbf{p}_g, \mathbf{q}))_{g \in \mathcal{G}} : \mathbf{q} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}) = \sigma \right\}$$

is piecewise linear (as the other components in the definition of \bar{m} are linear in π). This is the case since for each g , the mapping

$$\pi \in \Delta(\mathcal{G}) \longmapsto \left\{ \pi_g r(\mathbf{p}_g, \mathbf{q}) : \mathbf{q} \in \Delta(\mathcal{J}) \text{ such that } \tilde{H}(\mathbf{q}) = \sigma \right\}$$

is seen to be piecewise linear, by using the same one-dimensional argument as the one stated around (16) and also used in the proof of Corollary 31.

Step 3: We now exhibit the convergence rates. In view of the form of the defining set of constraints for \mathcal{C} , the coordinates of the elements in \mathcal{C} can be grouped according to each $g \in \mathcal{G}$ and projections onto \mathcal{C} can therefore be done separately for each such subset of coordinates. The subset of coordinates of $\sum_{t=1}^T \underline{r}(G_t, J_t)/T$ corresponding to a given g is formed by

$$\frac{N_T(g)}{T} r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \quad \text{and} \quad \frac{N_T(g)}{T} \tilde{H}(\hat{\mathbf{q}}_{T,g}).$$

When

$$\frac{N_T(g)}{T} r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \geq \max_{g' \in \mathcal{G}} \rho\left(\mathbf{p}_{g'}, \frac{N_T(g)}{T} \tilde{H}(\hat{\mathbf{q}}_{T,g})\right),$$

we denote these quantities by $\tilde{r}_{T,g}$ and $\tilde{\mathbf{v}}_{T,g}$. Otherwise, we project this pair on the set

$$\mathcal{C}_g = \left\{ (z_g, \mathbf{v}_g) \in \mathbb{R} \times \mathcal{F}_{\text{cone}} : z_g \geq \max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \mathbf{v}_g) \right\}$$

and denote by $\tilde{r}_{T,g}$ and $\tilde{\mathbf{v}}_{T,g}$ the coordinates of the projection; they satisfy the defining inequality of \mathcal{C}_g with equality,

$$\tilde{r}_{T,g} = \max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \tilde{\mathbf{v}}_{T,g}).$$

By distinguishing for each g according to which of the two cases above arose (for the first inequality), we may decompose and upper bound the swap regret as follows,

$$\begin{aligned} R_T^{\text{swap}} &= \sum_{g \in \mathcal{G}} \frac{N_T(g)}{T} \left(\max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \tilde{H}(\hat{\mathbf{q}}_{T,g})) - r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \right)_+ \\ &= \sum_{g \in \mathcal{G}} \left(\max_{g' \in \mathcal{G}} \rho\left(\mathbf{p}_{g'}, \frac{N_T(g)}{T} \tilde{H}(\hat{\mathbf{q}}_{T,g})\right) - \frac{N_T(g)}{T} r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \right)_+ \\ &\leq \sum_{g \in \mathcal{G}} \left| \max_{g' \in \mathcal{G}} \rho\left(\mathbf{p}_{g'}, \frac{N_T(g)}{T} \tilde{H}(\hat{\mathbf{q}}_{T,g})\right) - \max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \tilde{\mathbf{v}}_{g,T}) \right| + \sum_{g \in \mathcal{G}} \left| \tilde{r}_{T,g} - \frac{N_T(g)}{T} r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \right| \\ &\leq \sum_{g \in \mathcal{G}} \bar{L}_\rho \left\| \frac{N_T(g)}{T} \tilde{H}(\hat{\mathbf{q}}_{T,g}) - \tilde{\mathbf{v}}_{g,T} \right\|_2 + \sum_{g \in \mathcal{G}} \left| \tilde{r}_{T,g} - \frac{N_T(g)}{T} r(\mathbf{p}_g, \hat{\mathbf{q}}_{T,g}) \right|, \end{aligned}$$

where we used a fact proved in step 4, that the mapping

$$\mathbf{v} \in \mathcal{F}_{\text{cone}} \longmapsto \max_{g' \in \mathcal{G}} \rho(\mathbf{p}_{g'}, \mathbf{v}) \tag{18}$$

is \bar{L}_ρ -Lipschitz. In the last inequality we had a sum of ℓ^2 -norms, which can be bounded by a single ℓ^2 -norm,

$$\begin{aligned} R_T^{\text{swap}} &\leq \max\{\bar{L}_\rho, 1\} \sqrt{2N_{\mathcal{G}}} \left\| \left[\begin{array}{c} \tilde{r}_{T,g} \\ \tilde{\mathbf{v}}_{T,g} \end{array} \right]_{g \in \mathcal{G}} - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2 \\ &\leq \max\{\bar{L}_\rho, 1\} \sqrt{2N_{\mathcal{G}}} \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2, \end{aligned}$$

where we denoted by $N_{\mathcal{G}}$ the cardinality of \mathcal{G} . Resorting to the convergence rate stated in Theorem 26 concludes the proof.

Step 4: It only remains to prove the claimed Lipschitzness of the mapping (18). (All Lipschitzness statements that follow will be with respect to the ℓ^2 -norms.) To do so, it suffices to show that for all fixed elements $\mathbf{p} \in \Delta(\mathcal{I})$, the functions $\mathbf{v} \in \mathcal{F}_{\text{cone}} \mapsto \rho(\mathbf{p}, \mathbf{v})$ are Lipschitz, with a Lipschitz constant \bar{L}_ρ that is independent of \mathbf{p} . Note that we already proved at the end of the proof of Corollary 31 that $\sigma \in \mathcal{F} \mapsto \rho(\mathbf{p}, \sigma)$ is Lipschitz, with a Lipschitz constant L_ρ independent of \mathbf{p} . Consider now two elements $\mathbf{v}, \mathbf{v}' \in \mathcal{F}_{\text{cone}}$, which we write as $\mathbf{v} = \lambda \sigma$ and $\mathbf{v}' = \lambda' \sigma'$, with $\sigma, \sigma' \in \mathcal{F}$ and $\lambda, \lambda' \in \mathbb{R}_+$. Using triangle inequalities, the Lipschitzness of ρ on \mathcal{F} , and the fact that r thus ρ are bounded by R ,

$$\begin{aligned} |\rho(\mathbf{p}, \lambda \sigma) - \rho(\mathbf{p}, \lambda' \sigma')| &\leq |\lambda(\rho(\mathbf{p}, \sigma) - \rho(\mathbf{p}, \sigma'))| + |(\lambda - \lambda')\rho(\mathbf{p}, \sigma')| \\ &\leq \lambda L_\rho \|\sigma - \sigma'\|_2 + R |\lambda - \lambda'| \\ &\leq L_\rho \|\lambda \sigma - \lambda' \sigma' + (\lambda' - \lambda)\sigma'\|_2 + R |\lambda - \lambda'| \\ &\leq L_\rho \|\lambda \sigma - \lambda' \sigma'\|_2 + (R + L_\rho N_{\mathcal{I}}) |\lambda - \lambda'|, \end{aligned}$$

where we used also for the last inequality that since σ is a vector of $N_{\mathcal{I}}$ probability distributions over the signals, $\|\sigma\|_2 \leq \|\sigma\|_1 = N_{\mathcal{I}}$. To conclude the argument, we simply need to show that $|\lambda - \lambda'|$ can be bounded by $\|\lambda \sigma - \lambda' \sigma'\|_2$ up to some universal constant, which we do now. We resort again to the fact that $\|\sigma\|_1 = \|\sigma'\|_1 = N_{\mathcal{I}}$ and can thus write, thanks to a triangle inequality and assuming with no loss of generality that $\lambda' < \lambda$, that

$$|\lambda - \lambda'| = \frac{1}{N_{\mathcal{I}}} (\lambda \|\sigma\|_1 - \lambda' \|\sigma'\|_1) \leq \frac{1}{N_{\mathcal{I}}} \|\lambda \sigma - \lambda' \sigma'\|_1 \leq \frac{\sqrt{N_{\mathcal{H}} N_{\mathcal{I}}}}{N_{\mathcal{I}}} \|\lambda \sigma - \lambda' \sigma'\|_2,$$

where we used the Cauchy-Schwarz inequality for the final step. One can thus take, for instance,

$$\bar{L}_\rho = L_\rho + (R + L_\rho N_{\mathcal{I}}) \sqrt{\frac{N_{\mathcal{H}}}{N_{\mathcal{I}}}}.$$

This concludes the proof. ■

7. Summary of the Results

This paper extended Blackwell’s classical approachability theory to the case where set-valued functions are considered, which models ambiguity in the obtained reward. In the case of mixed actions taken, this extension was provided in the case of linear (Section 3) and concave–convex (Section 4) set-valued functions; only in the former case efficient strategies (up to a projection oracle) could be constructed.

The second part of this paper (Section 5) applies this theory of set-valued approachability to approachability with partial monitoring. The necessary and sufficient Condition (APM) for this was exhibited by Perchet (2011a) and was recalled in Section 5.1; its link with the necessary and sufficient condition for set-valued approachability was discussed in Section 5.2. Then, under a so-called assumption of bi-piecewise linearity of the game (r, H) at hand, an efficient strategy (up to a projection oracle) was constructed and studied in Section 5.3, for the approachability of any closed convex set \mathcal{C} . Alternatively, Section 5.4 showed that for any game (r, H) at hand but under the constraint that the target set \mathcal{C} is a polytope, the above efficient construction could still be used. In both cases, the novelty also relies not only the gained efficiency with respect to the construction by Perchet (2011a) but also on getting for the first time rates of convergence that are independent of the ambient dimension. The case of any game (r, H) and any closed target set \mathcal{C} was discussed, for the sake of completeness, at the end of Section 5.4, so that the present article contains a complete and self-contained constructive proof of the sufficiency of Condition (APM).

Finally, Section 6 showed that the well-studied case of regret minimization, a special case of approachability, could fall under the umbrella of bi-piecewise linearity, and hence be performed efficiently, as was already known.

Acknowledgments

Shie Mannor was partially supported by the ISF under contract 890015 and the Google Inter-university center for Electronic Markets and Auctions. Vianney Perchet benefited from the support of the ANR under grants ANR-10-BLAN 0112 and ANR-13-JS01-0004-01. Gilles Stoltz acknowledges support from Investissements d’Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047).

An extended abstract of this paper appeared in the *Proceedings of the 24th Annual Conference on Learning Theory (COLT’11)*, JMLR Workshop and Conference Proceedings, Volume 19, pages 515–536, 2011.

Appendix A. Proof of Theorem 15

Proof [of the second statement of Theorem 15] The proof of Corollary 7 extends to the case considered here and shows, thanks to the ad hoc consideration of the result stated in Lemma 6 as following from Definition 13, that for all $\mathbf{y} \in \Delta(\mathcal{B})$, the mapping $D_{\mathbf{y}}$ is still continuous over $\Delta(\mathcal{A})$. We now proceed by contradiction and assume that (SVAC) is not satisfied. The first part of the proof of the necessity of (SVAC) in Theorem 8 also applies to the present case: there exists \mathbf{y}_0 such that $D_{\mathbf{y}_0} \geq D_{\min} > 0$ over $\Delta(\mathcal{A})$. It then suffices to note that whenever the second player resorts to $\mathbf{y}_t = \mathbf{y}_0$ at all rounds $t \geq 1$, then for all

strategies of the first player, the quantity of interest in the set-valued approachability can be lower bounded as follows. Thanks to the concavity in the first argument,

$$\begin{aligned} & \sup \left\{ \inf_{c \in \mathcal{C}} \|\xi - c\|_2 : \xi \in \frac{1}{T} \sum_{t=1}^T \bar{m}(\mathbf{x}_t, \mathbf{y}_0) \right\} \\ & \geq \sup \left\{ \inf_{c \in \mathcal{C}} \|\xi - c\|_2 : \xi \in \bar{m} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t, \mathbf{y}_0 \right) \right\} = D_{\mathbf{y}_0} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \right) \geq D_{\min} > 0. \end{aligned}$$

Therefore, \mathcal{C} is \bar{m} -approachable by no strategy of the first player. ■

The proof of the first statement of Theorem 15 relies on the use of approximately calibrated strategies of the first player, as introduced and studied (among others) by Dawid (1982), Foster and Vohra (1998), Mannor and Stoltz (2010). Formally, given $\eta > 0$, an η -calibrated strategy of the first player considers some finite covering of $\Delta(\mathcal{B})$ by N_η balls of radius η and abides by the following constraints. Denoting by $\mathbf{y}^1, \dots, \mathbf{y}^{N_\eta}$ the centers of the balls in the covering (they form what will be referred to later on as an η -grid), such a strategy chooses only forecasts in $\{\mathbf{y}^1, \dots, \mathbf{y}^{N_\eta}\}$. We thus denote by L_t the index chosen in $\{1, \dots, N_\eta\}$ at round t and by

$$N_T(\ell) = \sum_{t=1}^T \mathbb{I}_{\{L_t = \ell\}}$$

the total number of rounds within the first T ones when the element ℓ of the grid was chosen. We denote by $(\cdot)_+$ the function that gives the nonnegative part of a real number. The final condition to be satisfied is that for all $\delta \in (0, 1)$, there exists an integer T_δ such that for all strategies of the second player, with probability at least $1 - \delta$, for all $T \geq T_\delta$,

$$\sum_{\ell=1}^{N_\eta} \frac{N_T(\ell)}{T} \left(\left\| \mathbf{y}^\ell - \frac{1}{N_T(\ell)} \sum_{t=1}^T \mathbf{y}_t \mathbb{I}_{\{L_t = \ell\}} \right\|_1 - \eta \right)_+ \leq \delta. \tag{19}$$

This calibration criterion is slightly stronger than the classical η -calibration score usually considered in the literature, which consists of omitting nonnegative parts in the criterion above and ensuring that for all strategies of the second player, with probability at least $1 - \delta$, for all $T \geq T_\delta$,

$$\sum_{\ell=1}^{N_\eta} \frac{N_T(\ell)}{T} \left\| \mathbf{y}^\ell - \frac{1}{N_T(\ell)} \sum_{t=1}^T \mathbf{y}_t \mathbb{I}_{\{L_t = \ell\}} \right\|_1 \leq \eta + \delta. \tag{20}$$

The existence of a calibrated strategy in the sense of (19) however follows from the same approachability-based construction studied in Mannor and Stoltz (2010) to get (20) and is detailed below in Section B. In the sequel we will only use the following consequence of calibration: that for all strategies of the second player, with probability at least $1 - \delta$, for all $T \geq T_\delta$,

$$\max_{\ell=1, \dots, N_\eta} \frac{N_T(\ell)}{T} \left(\left\| \mathbf{y}^\ell - \frac{1}{N_T(\ell)} \sum_{t=1}^T \mathbf{y}_t \mathbb{I}_{\{L_t = \ell\}} \right\|_1 - \eta \right)_+ \leq \delta. \tag{21}$$

Proof [of the first statement of Theorem 15] The insight of this proof is similar to the one illustrated in Perchet (2009). We first note that it suffices to prove that for all $\varepsilon > 0$, the set \mathcal{C}_ε defined as the ε -neighborhood of \mathcal{C} is \bar{m} -approachable. This is so up to proceeding in regimes $r = 1, 2, \dots$ each corresponding to a dyadic value $\varepsilon_r = 2^{-r}$ and lasting for a number of rounds carefully chosen in terms of the length of the previous regimes.

Therefore, we fix $\varepsilon > 0$ and associate with it a modulus of continuity $\eta > 0$ given by the uniform continuity of \bar{m} in its second argument. We consider an $\eta/2$ -calibrated strategy of the first player, which we will use as an auxiliary strategy. Since (SVAC) is satisfied, we may associate with each element \mathbf{y}^ℓ of the underlying $\eta/2$ -grid a mixed action $\mathbf{x}^\ell \in \Delta(\mathcal{A})$ such that $\bar{m}(\mathbf{x}^\ell, \mathbf{y}^\ell) \subseteq \mathcal{C}$. The main strategy of the first player then prescribes the use of $\mathbf{x}_t = \mathbf{x}^{L_t}$ at each round $t \geq 1$. The intuition behind this definition is that if \mathbf{y}^{L_t} is forecasted by the auxiliary strategy, then since the latter is calibrated, one should play as good as possible against \mathbf{y}^{L_t} . In view of the aim at hand, which is approaching \mathcal{C} , such a good reply is given by \mathbf{x}^{L_t} .

To assess the constructed strategy, we group rounds according to the values ℓ taken by the L_t . To that end, we recall that $N_T(\ell)$ denotes the number of rounds in which \mathbf{y}^ℓ was forecasted and \mathbf{x}^ℓ was played. The average payoff up to round T is then rewritten as

$$\frac{1}{T} \sum_{t=1}^T \bar{m}(\mathbf{x}_t, \mathbf{y}_t) = \sum_{\ell=1}^{N_{\eta/2}} \frac{N_T(\ell)}{T} \left(\frac{1}{N_T(\ell)} \sum_{t=1}^T \bar{m}(\mathbf{x}^\ell, \mathbf{y}_t) \mathbb{I}_{\{L_t=\ell\}} \right).$$

We denote for all ℓ such that $N_T(\ell) > 0$ the average of their corresponding mixed actions \mathbf{y}_t by

$$\bar{\mathbf{y}}_T^\ell = \frac{1}{N_T(\ell)} \sum_{t=1}^T \mathbf{y}_t \mathbb{I}_{\{L_t=\ell\}}.$$

The convexity of \bar{m} in its second argument leads to the inclusion

$$\frac{1}{T} \sum_{t=1}^T \bar{m}(\mathbf{x}_t, \mathbf{y}_t) = \sum_{\ell=1}^{N_{\eta/2}} \frac{N_T(\ell)}{T} \left(\frac{1}{N_T(\ell)} \sum_{t=1}^T \bar{m}(\mathbf{x}^\ell, \mathbf{y}_t) \mathbb{I}_{\{L_t=\ell\}} \right) \subseteq \sum_{\ell=1}^{N_{\eta/2}} \frac{N_T(\ell)}{T} \bar{m}(\mathbf{x}^\ell, \bar{\mathbf{y}}_T^\ell).$$

We recall that \mathbf{B} denotes the unit Euclidean ball in \mathbb{R}^d . To show that the above-defined strategy \bar{m} -approaches $\mathcal{C}_\varepsilon = \mathcal{C} + \varepsilon \mathbf{B}$, it suffices to show that for all $\delta \in (0, 1)$, there exists an integer T'_δ such that for all strategies of the second player,

$$\mathbb{P} \left\{ \forall T \geq T'_\delta, \quad \sum_{\ell=1}^{N_{\eta/2}} \frac{N_T(\ell)}{T} \bar{m}(\mathbf{x}^\ell, \bar{\mathbf{y}}_T^\ell) \subseteq \mathcal{C} + (\varepsilon + \delta) \mathbf{B} \right\} \geq 1 - \delta.$$

We denote by M a bound in ℓ^2 -norm on \bar{m} , i.e., for all $\mathbf{x} \in \Delta(\mathcal{A})$ and $\mathbf{y} \in \Delta(\mathcal{B})$, the inclusion $\bar{m}(\mathbf{x}, \mathbf{y}) \subseteq M \mathbf{B}$ holds. We let $\delta' = \delta(\eta/2)/(M N_{\eta/2})$ and define T'_δ as the time $T_{\delta'}$ corresponding to (21). All statements that follow will be for all strategies of the second player and with probability at least $1 - \delta' \geq 1 - \delta$, for all $T \geq T'_\delta$, as required. For each index

ℓ of the grid, either $\delta'T/N_T(\ell) \leq \eta/2$ or $\delta'T/N_T(\ell) > \eta/2$. In the first case, following (21), $\|\mathbf{y}^\ell - \bar{\mathbf{y}}_T^\ell\| \leq \eta/2 + \delta'T/N_T(\ell) \leq \eta$; since η is the modulus of continuity for ε , we get that

$$\frac{N_T(\ell)}{T} \bar{m}(\mathbf{x}^\ell, \bar{\mathbf{y}}_T^\ell) \subseteq \frac{N_T(\ell)}{T} \left(\bar{m}(\mathbf{x}^\ell, \mathbf{y}^\ell) + \varepsilon \mathbf{B} \right) \subseteq \frac{N_T(\ell)}{T} (\mathcal{C} + \varepsilon \mathbf{B}),$$

where we used the definition of \mathbf{x}^ℓ to get the second inclusion. In the second case, using the boundedness of \bar{m} , we simply write

$$\frac{N_T(\ell)}{T} \bar{m}(\mathbf{x}^\ell, \bar{\mathbf{y}}_T^\ell) \subseteq \frac{N_T(\ell)}{T} M \mathbf{B} \subseteq \frac{\delta'}{\eta/2} M \mathbf{B}.$$

Summing these bounds over ℓ yields

$$\sum_{\ell=1}^{N_{\eta/2}} \frac{N_T(\ell)}{T} \bar{m}(\mathbf{x}^\ell, \bar{\mathbf{y}}_T^\ell) \subseteq \mathcal{C} + \varepsilon \mathbf{B} + \frac{N_{\eta/2} \delta'}{\eta/2} M \mathbf{B} = \mathcal{C} + (\varepsilon + \delta) \mathbf{B},$$

where we used the definition of δ' in terms of δ . This concludes the proof. ■

Appendix B. An Auxiliary Result of Calibration

We prove here (19) for a given $\eta > 0$ and do so by following closely the methodology of Mannor and Stoltz (2010). (Note that this result is of independent interest.)

We actually assume that the covering $\mathbf{y}^1, \dots, \mathbf{y}^{N_\eta}$ is slightly finer than what was required around (19) and that it forms an $\eta/N_{\mathcal{B}}$ -grid of $\Delta(\mathcal{B})$, i.e., that for all $\mathbf{y} \in \Delta(\mathcal{B})$, there exists $\ell \in \{1, \dots, N_\eta\}$ such that $\|\mathbf{y} - \mathbf{y}^\ell\|_1 \leq \eta/N_{\mathcal{B}}$.

We recall that elements $\mathbf{y} \in \mathcal{B}$ are denoted by $\mathbf{y} = (y_b)_{b \in \mathcal{B}}$ and we identify $\Delta(\mathcal{B})$ with a subset of $\mathbb{R}^{N_{\mathcal{B}}}$. In particular, \mathbb{I}_b , the Dirac mass on a given $b \in \mathcal{B}$, is a binary vector whose only non-null component is the one indexed by b . Finally, we denote by

$$\underline{0} = (0, \dots, 0) \quad \text{and} \quad \underline{1} = (1, \dots, 1)$$

the elements of $\mathbb{R}^{\mathcal{B}}$ respectively formed by zeros and ones only.

We consider a vector-valued payoff function $C : \{1, \dots, N_\eta\} \times \mathcal{B} \rightarrow \mathbb{R}^{2N_\eta N_{\mathcal{B}}}$ defined as follows; for all $\ell \in \{1, \dots, N_\eta\}$ and for all $b \in \mathcal{B}$,

$$C(\ell, b) = \left(\underline{0}, \dots, \underline{0}, \mathbf{y}^\ell - \mathbb{I}_b - \frac{\eta}{N_{\mathcal{B}}} \underline{1}, \mathbb{I}_b - \mathbf{y}^\ell - \frac{\eta}{N_{\mathcal{B}}} \underline{1}, \underline{0}, \dots, \underline{0} \right),$$

which is a vector of $2N_\eta$ elements of $\mathbb{R}^{\mathcal{B}}$ composed by $2(N_\eta - 1)$ occurrences of the zero element $\underline{0} \in \mathbb{R}^{\mathcal{B}}$ and two non-zero elements, located in the positions indexed by $2\ell - 1$ and 2ℓ .

We now show that the closed convex set $(\mathbb{R}_-)^{2N_\eta N_{\mathcal{B}}}$ is C -approachable; to do so, we resort to the characterization stated in Theorem 2. To each $\mathbf{y} \in \Delta(\mathcal{B})$ we will associate a pure action $\ell_{\mathbf{y}}$ in $\{1, \dots, N_\eta\}$ so that $C(\ell_{\mathbf{y}}, \mathbf{y}) \in (\mathbb{R}_-)^{2N_\eta N_{\mathcal{B}}}$; note that to satisfy the necessary and sufficient condition, it is not necessary here to resort to mixed actions of the

first player. The index $\ell_{\mathbf{y}}$ is any index ℓ such that $\|\mathbf{y} - \mathbf{y}^\ell\|_1 \leq \eta/N_{\mathcal{B}}$; such an index always exists as noted at the beginning of this proof. Indeed, one then has in particular that for each component $b \in \mathcal{B}$,

$$|y_b^{\ell_{\mathbf{y}}} - y_b| \leq \|\mathbf{y}^{\ell_{\mathbf{y}}} - \mathbf{y}\|_1 \leq \eta/N_{\mathcal{B}}.$$

A straightforward adaptation of the proof of Theorem 3 (see, e.g., Mertens et al., 1994) then yields a strategy such that for all $\delta \in (0, 1)$ and for all strategies of the second player, with probability at least $1 - \delta$,

$$\sup_{\tau \geq T} \inf_{c \in (\mathbb{R}_-)^{2N_{\eta}N_{\mathcal{B}}}} \left\| c - \frac{1}{\tau} \sum_{t=1}^{\tau} C(L_t, \mathbf{y}_t) \right\|_2 \leq 2M \sqrt{\frac{2}{\delta T}}, \tag{22}$$

where M is a bound in Euclidean norm over C , e.g., $M = 4 + 2\eta$. The quantities of interest can be rewritten as

$$\frac{1}{\tau} \sum_{t=1}^{\tau} C(L_t, \mathbf{y}_t) = \left(\frac{N_{\tau}(\ell)}{\tau} (\mathbf{y}^\ell - \bar{\mathbf{y}}_{\tau}^{\ell}) - \frac{N_{\tau}(\ell)}{\tau} \frac{\eta}{N_{\mathcal{B}}} \mathbf{1}, \frac{N_{\tau}(\ell)}{\tau} (\bar{\mathbf{y}}_{\tau}^{\ell} - \mathbf{y}^\ell) - \frac{N_{\tau}(\ell)}{\tau} \frac{\eta}{N_{\mathcal{B}}} \mathbf{1} \right)_{\ell \in \{1, \dots, N_{\eta}\}},$$

where we recall that we denoted for all ℓ such that $N_{\tau}(\ell) > 0$ the average of their corresponding mixed actions \mathbf{y}_t by

$$\bar{\mathbf{y}}_{\tau}^{\ell} = \frac{1}{N_{\tau}(\ell)} \sum_{t=1}^{\tau} \mathbf{y}_t \mathbb{I}_{\{L_t = \ell\}}.$$

The projection in ℓ^2 -norm of quantity of interest onto $(\mathbb{R}_-)^{2N_{\eta}N_{\mathcal{B}}}$ is formed by its non-positive components, so that its square distance to $(\mathbb{R}_-)^{2N_{\eta}N_{\mathcal{B}}}$ equals

$$\begin{aligned} \inf_{c \in (\mathbb{R}_-)^{2N_{\eta}N_{\mathcal{B}}}} \left\| c - \frac{1}{\tau} \sum_{t=1}^{\tau} C(L_t, \mathbf{y}_t) \right\|_2^2 &= \sum_{\ell=1}^{N_{\eta}} \left(\frac{N_{\tau}(\ell)}{\tau} \right)^2 \sum_{b \in \mathcal{B}} \underbrace{\left(\left(y_b^{\ell} - \bar{y}_{\tau, b}^{\ell} - \frac{\eta}{N_{\mathcal{B}}} \right)_+^2 + \left(\bar{y}_{\tau, b}^{\ell} - y_b^{\ell} - \frac{\eta}{N_{\mathcal{B}}} \right)_+^2 \right)}_{= (|\bar{y}_{\tau, b}^{\ell} - y_b^{\ell}| - \eta/N_{\mathcal{B}})_+^2}. \end{aligned}$$

Therefore, our target is achieved: using the fact that $(\cdot)_+$ is subadditive first, and then applying the Cauchy-Schwarz inequality,

$$\begin{aligned} \sum_{\ell=1}^{N_{\eta}} \frac{N_{\tau}(\ell)}{\tau} \left(\|\mathbf{y}^{\ell} - \bar{\mathbf{y}}_{\tau}\|_1 - \eta \right)_+ &\leq \sum_{\ell=1}^{N_{\eta}} \frac{N_{\tau}(\ell)}{\tau} \sum_{b \in \mathcal{B}} \left(|y_b^{\ell} - \bar{y}_{\tau, b}^{\ell}| - \frac{\eta}{N_{\mathcal{B}}} \right)_+ \\ &\leq \sqrt{N_{\eta}N_{\mathcal{B}}} \sqrt{\sum_{\ell=1}^{N_{\eta}} \left(\frac{N_{\tau}(\ell)}{\tau} \right)^2 \sum_{b \in \mathcal{B}} \left(|y_b^{\ell} - \bar{y}_{\tau, b}^{\ell}| - \frac{\eta}{N_{\mathcal{B}}} \right)_+^2} \\ &\leq 2M \sqrt{N_{\eta}N_{\mathcal{B}}} \sqrt{\frac{2}{\delta T}}, \end{aligned}$$

where the last inequality holds, by (22), for all $\tau \geq T$ with probability at least $1 - \delta$. Choosing an integer T_δ sufficiently large so that

$$2M\sqrt{N_\eta N_\mathcal{B}} \sqrt{\frac{2}{\delta T_\delta}} \leq \delta$$

concludes the proof of the property stated in (19).

Appendix C. Proof of Lemma 23

Proof For all $(i, j) \in \mathcal{I} \times \mathcal{J}$, the quantity $H(i, j)$ is a probability distribution over the set of signals \mathcal{H} ; we denote by $H_s(i, j)$ the probability mass that it puts on some signal $s \in \mathcal{H}$.

Equation (7) indicates that for each pair $(i, s) \in \mathcal{I} \times \mathcal{H}$,

$$\sum_{t=(n-1)L+1}^{nL} \left(\frac{\mathbb{I}_{\{S_t=s\}}\mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} - H_s(i, J_t) \right)$$

is a sum of L elements of a martingale difference sequence, with respect to the filtration whose t -th element is generated by \mathbf{p}_n , the pairs (I_s, S_s) for $s \leq t$, and J_s for $s \leq t + 1$. The conditional variances of the increments are bounded by

$$\mathbb{E}_t \left[\left(\frac{\mathbb{I}_{\{S_t=s\}}\mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} \right)^2 \right] \leq \frac{1}{p_{i,n}^2} \mathbb{E}_t [\mathbb{I}_{\{I_t=i\}}] = \frac{1}{p_{i,n}};$$

since by definition of the strategy, $\mathbf{p}_n = (1 - \gamma) \mathbf{x}_n + \gamma \mathbf{u}$, we have that $p_{i,n} \geq \gamma/N_{\mathcal{I}}$, which shows that the sum of the conditional variances is bounded by

$$\sum_{t=(n-1)L+1}^{nL} \text{Var}_t \left(\frac{\mathbb{I}_{\{S_t=s\}}\mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} \right) \leq \frac{LN_{\mathcal{I}}}{\gamma}.$$

The Bernstein-Freedman inequality (see Freedman, 1975 or Cesa-Bianchi et al., 2006, Lemma A.1) therefore indicates that with probability at least $1 - \delta$,

$$\left| \frac{1}{L} \sum_{t=(n-1)L+1}^{nL} \frac{\mathbb{I}_{\{S_t=s\}}\mathbb{I}_{\{I_t=i\}}}{p_{I_t,n}} - \underbrace{\frac{1}{L} \sum_{t=(n-1)L+1}^{nL} H_s(i, J_t)}_{= H_s(i, \hat{\mathbf{q}}_n)} \right| \leq \sqrt{2 \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2}{\delta}.$$

Therefore, by summing the above inequalities over $i \in \mathcal{I}$ and $s \in \mathcal{H}$, we get (after a union bound) that with probability at least $1 - N_{\mathcal{I}}N_{\mathcal{H}}\delta$,

$$\left\| \tilde{\sigma}_n - \tilde{H}(\hat{\mathbf{q}}_n) \right\|_2 \leq \sqrt{N_{\mathcal{I}}N_{\mathcal{H}}} \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma L} \ln \frac{2}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma L} \ln \frac{2}{\delta} \right).$$

Finally, since $\hat{\sigma}_n$ is the projection in the ℓ^2 -norm of $\tilde{\sigma}_n$ onto the convex set \mathcal{F} , to which $\tilde{H}(\hat{\mathbf{q}}_n)$ belongs, we have that

$$\left\| \hat{\sigma}_n - \tilde{H}(\hat{\mathbf{q}}_n) \right\|_2 \leq \left\| \tilde{\sigma}_n - \tilde{H}(\hat{\mathbf{q}}_n) \right\|_2,$$

and this concludes the proof. ■

Appendix D. Proof of Theorem 25

Proof We denote by d_t^α the squared distance of \widehat{m}_t^α to \mathcal{C} ,

$$d_t^\alpha = \inf_{c \in \mathcal{C}} \|c - \widehat{m}_t^\alpha\|^2 = \|c_t^\alpha - \widehat{m}_t^\alpha\|^2$$

and use the shortcut notation $m_t = m(\mathbf{x}_t, \mathbf{y}_t)$ for all $t \geq 1$. Then,

$$\begin{aligned} d_{t+1}^\alpha &\leq \|\widehat{m}_{t+1}^\alpha - c_t^\alpha\|^2 = \left\| \widehat{m}_t^\alpha - c_t^\alpha + \frac{(t+1)^\alpha}{T_{t+1}^\alpha} (m_{t+1} - \widehat{m}_t^\alpha) \right\|^2 \\ &\leq \|\widehat{m}_t^\alpha - c_t^\alpha\|^2 + \frac{2(t+1)^\alpha}{T_{t+1}^\alpha} \langle \widehat{m}_t^\alpha - c_t^\alpha, m_{t+1} - m_t^\alpha \rangle + \left(\frac{(t+1)^\alpha}{T_{t+1}^\alpha} \right)^2 \|m_{t+1} - \widehat{m}_t^\alpha\|^2 \\ &\leq d_t^\alpha + \frac{2(t+1)^\alpha}{T_{t+1}^\alpha} \left(\underbrace{\langle \widehat{m}_t^\alpha - c_t^\alpha, m_{t+1} - c_t^\alpha \rangle}_{\leq 0} + \langle \widehat{m}_t^\alpha - c_t^\alpha, c_t^\alpha - m_t^\alpha \rangle \right) + \left(\frac{(t+1)^\alpha}{T_{t+1}^\alpha} \right)^2 4M^2 \\ &\leq d_t^\alpha \left(1 - \frac{2(t+1)^\alpha}{T_{t+1}^\alpha} \right) + \left(\frac{(t+1)^\alpha}{T_{t+1}^\alpha} \right)^2 4M^2, \end{aligned}$$

where we used in the third inequality the same convex projection inequality as in the proof of Theorem 3.

The first inequality in (15) then follows by induction: the bound $2M$ for $t = 1$ is by boundedness of m . If the stated bound holds for d_t^α , then

$$d_{t+1}^\alpha \leq \left(2M \frac{\sqrt{\sum_{s=1}^t s^{2\alpha}}}{\sum_{s=1}^t s^\alpha} \right)^2 \left(1 - \frac{2(t+1)^\alpha}{T_{t+1}^\alpha} \right) + \left(\frac{(t+1)^\alpha}{T_{t+1}^\alpha} \right)^2 4M^2 \leq 4M^2 \frac{\sum_{s=1}^{t+1} s^{2\alpha}}{(T_{t+1}^\alpha)^2},$$

as desired, since

$$\frac{1}{(T_t^\alpha)^2} \left(1 - \frac{2(t+1)^\alpha}{T_{t+1}^\alpha} \right) = \frac{T_t^\alpha - (t+1)^\alpha}{T_{t+1}^\alpha (T_t^\alpha)^2} = \frac{1}{T_{t+1}^\alpha (T_t^\alpha)^2} \frac{(T_t^\alpha)^2 - (t+1)^{2\alpha}}{T_t^\alpha + (t+1)^\alpha} \leq \frac{1}{(T_{t+1}^\alpha)^2}.$$

The second inequality in (15) is straightforward for $\alpha = 0$ and is proved for $\alpha > 0$ as follows. First, by comparing sums and integrals, we get that for all $t \geq 1$,

$$\frac{t^{\alpha+1}}{\alpha+1} = \int_0^t s^\alpha ds \leq \sum_{s=1}^t s^\alpha \leq t \times t^\alpha = t^{\alpha+1}.$$

Therefore,

$$\frac{\sqrt{\sum_{s=1}^t s^{2\alpha}}}{\sum_{s=1}^t s^\alpha} \leq (\alpha+1) \frac{\sqrt{t^{2\alpha+1}}}{t^{\alpha+1}} = \frac{\alpha+1}{\sqrt{t}}.$$

This concludes the proof. Note for later purposes that upper bounding above the sum of the s^α as

$$\sum_{s=1}^t s^\alpha \leq t^\alpha + \int_1^t s^\alpha ds \leq t^\alpha + \frac{t^{\alpha+1}}{\alpha+1}$$

shows that

$$\sum_{s=1}^t s^\alpha \sim \frac{t^{\alpha+1}}{\alpha+1}. \quad \blacksquare$$

Appendix E. Proof of Theorem 26

Proof The proof follows closely the proof of Theorem 24. We choose N so as to write $T = T_N^\alpha + k$ where $0 \leq k \leq L_{N+1} - 1$. We adapt step 1 as follows,

$$\left\| \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) - \frac{1}{T_N^\alpha} \sum_{t=1}^{T_N^\alpha} r(I_t, J_t) \right\|_2 \leq R \left(\frac{k}{T} + \left(\frac{1}{T_N^\alpha} - \frac{1}{T} \right) T_N^\alpha \right) = \frac{2k}{T} R \leq \frac{2L_{N+1}}{T} R.$$

Second, as in step 2, we resort again to the Hoeffding-Azuma inequality for sums of Hilbert space-valued martingale differences; with probability at least $1 - \delta$,

$$\left\| \frac{1}{T_N^\alpha} \sum_{t=1}^{T_N^\alpha} r(I_t, J_t) - \frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha r(\mathbf{p}_n, \hat{\mathbf{q}}_n) \right\|_2 \leq 4R \sqrt{\frac{\ln(2/\delta)}{T_N^\alpha}}.$$

In view of the choice $\gamma_n = n^{-\alpha/3}$, step 3 translates here to

$$\begin{aligned} \left\| \frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha r(\mathbf{p}_n, \hat{\mathbf{q}}_n) - \frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha r(\mathbf{x}_n, \hat{\mathbf{q}}_n) \right\|_2 \\ \leq 2R \frac{\sum_{n=1}^N n^\alpha \gamma_n}{T_N^\alpha} = 2R \frac{\sum_{n=1}^N n^{2\alpha/3}}{T_N^\alpha} = 2R \frac{T_N^{(2\alpha/3)}}{T_N^\alpha}. \end{aligned}$$

The same argument as the one at the beginning of the proof of Theorem 24 shows that

$$\frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha r(\mathbf{x}_n, \hat{\mathbf{q}}_n) \in \frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha \overline{\mathbf{m}} \left(\boldsymbol{\theta}_n, \Phi \left(\tilde{H}(\hat{\mathbf{q}}_n) \right) \right).$$

Step 4 starts also by an application of Lemma 23 together with the Lipschitzness of Φ to get that for all regimes $n = 1, \dots, N$, with probability at least $1 - \delta$,

$$\left\| \Phi(\hat{\sigma}_n) - \Phi \left(\tilde{H}(\hat{\mathbf{q}}_n) \right) \right\|_2 \leq \kappa_\Phi \sqrt{N_{\mathcal{I}} N_{\mathcal{H}}} \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma_n L_n} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma_n L_n} \ln \frac{2N_{\mathcal{I}} N_{\mathcal{H}}}{\delta} \right).$$

By a union bound, the above bound holds for all regimes $n = 1, \dots, N$ with probability at least $1 - N\delta$. Then, an application of Lemma 6 shows that

$$\frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha \overline{\mathbf{m}} \left(\boldsymbol{\theta}_n, \Phi \left(\tilde{H}(\hat{\mathbf{q}}_n) \right) \right) \text{ is in a } \varepsilon_N\text{-neighborhood of } \frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha \overline{\mathbf{m}} \left(\boldsymbol{\theta}_n, \Phi(\hat{\sigma}_n) \right),$$

where, substituting the values of $L_n = n^\alpha$ and $\gamma_n = n^{-\alpha/3}$,

$$\begin{aligned} \varepsilon_N &= R\sqrt{N_{\mathcal{B}}} \left(\kappa_{\Phi} \sqrt{N_{\mathcal{I}}N_{\mathcal{H}}} \frac{1}{T_N^\alpha} \sum_{n=1}^N n^\alpha \left(\sqrt{\frac{2N_{\mathcal{I}}}{\gamma_n L_n} \ln \frac{2N_{\mathcal{I}}N_{\mathcal{H}}}{\delta}} + \frac{1}{3} \frac{N_{\mathcal{I}}}{\gamma_n L_n} \ln \frac{2N_{\mathcal{I}}N_{\mathcal{H}}}{\delta} \right) \right) \\ &= R\sqrt{N_{\mathcal{B}}} \left(\kappa_{\Phi} \sqrt{N_{\mathcal{I}}N_{\mathcal{H}}} \left(\frac{T_N^{(2\alpha/3)}}{T_N^\alpha} \sqrt{2N_{\mathcal{I}} \ln \frac{2N_{\mathcal{I}}N_{\mathcal{H}}}{\delta}} + \frac{T_N^{(\alpha/3)}}{T_N^\alpha} \frac{N_{\mathcal{I}}}{3} \ln \frac{2N_{\mathcal{I}}N_{\mathcal{H}}}{\delta} \right) \right). \end{aligned}$$

It then suffices, as in step 5 of the original proof, to write the convergence rates for set-valued approachability guaranteed by the strategy Ψ . By combining the result of Lemma 10 with Theorem 25 and Lemma 6, we get

$$\inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T_n^\alpha} \sum_{n=1}^N n^\alpha \bar{m}(\theta_n, \Phi(\hat{\sigma}_n)) \right\|_2 \leq \frac{2R(\alpha+1)}{\sqrt{N}} \sqrt{N_{\mathcal{A}}N_{\mathcal{B}}}.$$

Putting all things together and applying a union bound, we obtain that with probability at least $1 - (N+1)\delta$,

$$\begin{aligned} \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \right\|_2 \\ = O \left(\frac{(N+1)^\alpha}{T} + \sqrt{\frac{\ln(1/\delta)}{T_N^\alpha}} + \frac{T_N^{(2\alpha/3)}}{T_N^\alpha} + \frac{T_N^{(2\alpha/3)}}{T_N^\alpha} \sqrt{\ln \frac{1}{\delta}} + \frac{T_N^{(\alpha/3)}}{T_N^\alpha} \ln \frac{1}{\delta} + \frac{1}{\sqrt{N}} \right). \end{aligned}$$

Since (as proved at the end of the proof of Theorem 25) $T_N^\beta \sim N^{\beta+1}/(\beta+1)$ for all $\beta \geq 0$, we get that

$$N \sim ((\alpha+1)T)^{1/(\alpha+1)} \quad \text{and} \quad T_N^\beta \sim \frac{N^{\beta+1}}{\beta+1} \sim \kappa_{\alpha,\beta} T^{(\beta+1)/(\alpha+1)},$$

where $\kappa_{\alpha,\beta}$ is a constant that only depends on α and β . Replacing δ by $\delta/(N+1)$ as we did in step 5 of the proof of Theorem 24, choosing $\alpha = 3/2$ and substituting the equivalences above ensures the result. ■

References

- J. Abernethy, P. L. Bartlett, and E. Hazan. Blackwell approachability and low-regret learning are equivalent. In *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*. JMLR Workshop and Conference Proceedings, 2011.
- G. Bartók, D. Pál, and C. Szepesvári. Toward a classification of finite partial-monitoring games. In *Proceedings of the Twenty-First International Conference on Algorithmic Learning Theory (ALT'10)*, pages 224–238. Springer, 2010.
- G. Bartók, D. Pál, and C. Szepesvári. Minimax regret of finite partial-monitoring games in stochastic environments. In *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*. JMLR Workshop and Conference Proceedings, 2011.

- D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6:1–8, 1956a.
- D. Blackwell. Controlled random walks. In *Proceedings of the International Congress of Mathematicians, 1954, Amsterdam, vol. III*, pages 336–338, 1956b.
- A. Blum and Y. Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, 2007.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- N. Cesa-Bianchi, G. Lugosi, and G. Stoltz. Regret minimization under partial monitoring. *Mathematics of Operations Research*, 31:562–580, 2006.
- X. Chen and H. White. Laws of large numbers for Hilbert space-valued mixingales with applications. *Econometric Theory*, 12:284–304, 1996.
- A.P. Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77:605–613, 1982.
- D. Foster and R. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.
- D. Foster and R. Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, 29:7–36, 1999.
- D.P. Foster and A. Rakhlin. No internal regret via neighborhood watch. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics (AISTATS’12)*, 2012.
- D.A. Freedman. On tail probabilities for martingales. *Annals of Probability*, 3:100–118, 1975.
- J.E. Goodman and J. O’Rourke, editors. *Handbook of Discrete and Computational Geometry*. Discrete Mathematics and its Applications. Chapman & Hall/CRC, Boca Raton, FL, second edition, 2004.
- 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*, 98:26–54, 2001.
- E. Lehrer and E. Solan. Learning to play partially-specified equilibrium. Mimeo, 2007.
- G. Lugosi, S. Mannor, and G. Stoltz. Strategies for prediction under imperfect monitoring. *Mathematics of Operations Research*, 33:513–528, 2008. An extended abstract was presented at COLT’07.
- S. Mannor and N. Shimkin. On-line learning with imperfect monitoring. In *Proceedings of the Sixteenth Annual Conference on Learning Theory (COLT’03)*, pages 552–567. Springer, 2003.

- 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 and G. Stoltz. A geometric proof of calibration. *Mathematics of Operations Research*, 35:721–727, 2010.
- S. Mannor, J. Tsitsiklis, and J. Y. Yu. Online learning with sample path constraints. *Journal of Machine Learning Research*, 10(Mar):569–590, 2009.
- J.-F. Mertens, S. Sorin, and S. Zamir. Repeated games. Technical Report no. 9420, 9421, 9422, Université de Louvain-la-Neuve, 1994.
- V. Perchet. Calibration and internal no-regret with random signals. In *Proceedings of the Twentieth International Conference on Algorithmic Learning Theory (ALT'09)*, pages 68–82, 2009.
- V. Perchet. Approachability of convex sets in games with partial monitoring. *Journal of Optimization Theory and Applications*, 149:665–677, 2011a.
- V. Perchet. Internal regret with partial monitoring calibration-based optimal algorithms. *Journal of Machine Learning Research*, 12:1893–1921, 2011b.
- V. Perchet and M. Quincampoix. On an unified framework for approachability in games with or without signals. Mimeo, 2011.
- A. Piccolboni and C. Schindelhauer. Discrete prediction games with arbitrary feedback and loss. In *Proceedings of the Fourteenth Annual Conference on Computational Learning Theory (COLT'01)*, pages 208–223, 2001.
- A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Beyond regret. In *Proceedings of the Twenty-Fourth Annual Conference on Learning Theory (COLT'11)*. JMLR Workshop and Conference Proceedings, 2011.
- J. Rambau and G. Ziegler. Projections of polytopes and the generalized Baues conjecture. *Discrete and Computational Geometry*, 16:215–237, 1996.
- A. Rustichini. Minimizing regret: The general case. *Games and Economic Behavior*, 29: 224–243, 1999.