



HAL
open science

Corrigendum to "Robust approachability and regret minimization in games with partial monitoring"

Shie Mannor, Vianney Perchet, Gilles Stoltz

► **To cite this version:**

Shie Mannor, Vianney Perchet, Gilles Stoltz. Corrigendum to "Robust approachability and regret minimization in games with partial monitoring". 2011. hal-00617554

HAL Id: hal-00617554

<https://hal.science/hal-00617554>

Preprint submitted on 29 Aug 2011

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

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

Corrigendum to “Robust Approachability and Regret Minimization in Games with Partial Monitoring”

July 7, 2011

Shie Mannor

Technion, Haifa
Israel

shie@ee.technion.ac.il

Vianney Perchet

Ecole normale supérieure, Cachan
France

vianney.perchet@normalesup.org

Gilles Stoltz

Ecole normale supérieure,*Paris
& HEC Paris

France
gilles.stoltz@ens.fr

In the analyses performed in Section 4 of the original paper, we used a linearity property which we realized was incorrect; all mentioned results still hold up to a slight twist on the algorithms to be used. We use below the notation of the original paper without explicitly redefining them.

1 Description of Error

The incorrect statement can be found at the end of the proof of the geometric lemma (Lemma 10):

[...] The proof is concluded by noting that by definition, for all $\sigma \in \mathcal{F}$, the applications $\mathbf{p} \in \Delta(\mathcal{I}) \mapsto \overline{m}(\mathbf{p}, \sigma)$ are linear.

A similar incorrect property of linearity in the first argument was also used in the (warm-up) Section 4.1. However, with the needed correction, the special case of Section 4.1 will no longer be much easier to handle than the general result; hence, this section should simply be discarded.

The example below illustrates why \overline{m} is in general not linear in its first argument (just as it is not linear in its second argument neither).

Example 1 We consider a game in which the second player (when playing L and M) can force the first player to play a game of matching pennies in the dark; in the matrix below, the real numbers denote the payoff while \clubsuit and \heartsuit denote the two possible signals.

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

A straightforward calculation shows that

$$\overline{m}(\delta_T, \clubsuit) = \overline{m}(\delta_B, \clubsuit) = [-1, 1] \quad \text{while} \quad \overline{m}\left(\frac{1}{2}\delta_T + \frac{1}{2}\delta_B, \clubsuit\right) = \{0\}.$$

Actually, the inclusion

$$\overline{m}(\mathbf{p}, \sigma) \subseteq \sum_{i \in \mathcal{I}} p_i \overline{m}(i, \sigma)$$

is true for all games, all $\mathbf{p} \in \Delta(\mathcal{I})$, and $\sigma \in \mathcal{F}$, but it is a strict inclusion in general. Therefore, the linear extension $\overline{\overline{m}}$ to $\Delta(\mathcal{I} \times \mathcal{B})$ of the restriction of \overline{m} to $\mathcal{I} \times \mathcal{B}$ considered in the original paper does not coincide with \overline{m} , contrary to what is stated in Definition 11. Yet, we need some linear set-valued mapping to apply the general results of Section 3.

*CNRS – Ecole normale supérieure, Paris – INRIA, within the project-team CLASSIC

A study of the properties of the mapping $\mathbf{p} \mapsto \overline{m}(\mathbf{p}, \clubsuit)$ sheds however some light on a possible patch, along the same lines as the construction already exhibited in Lemma 10 to get some linearity for \overline{m} in its second argument, up to the extension of the set of pure actions \mathcal{I} to a finite set of possibly mixed actions \mathcal{B} .

Indeed, $\mathbf{p} \mapsto \overline{m}(\mathbf{p}, \clubsuit)$ is piecewise linear, which can be seen by introducing a set \mathcal{A} of possibly mixed actions extending the set $\mathcal{I} = \{T, B\}$ of pure actions and containing

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

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 $\lambda \geq 1/2$, by definition of \overline{m} ,

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

so that we have the convex decomposition

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

which can be restated as

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

That is, $\overline{m}(\cdot, \clubsuit)$ is linear on the subset of $\Delta(\mathcal{I})$ corresponding to mixed actions \mathbf{p}_λ with $\lambda \geq 1/2$. Since a similar property holds the subset of distributions with $\lambda \leq 1/2$, we have proved that $\overline{m}(\cdot, \clubsuit)$ is piecewise linear on $\Delta(\mathcal{I})$.

We now prove that —just as in Lemma 10 of the original paper— this entails a linearity property on a lifted space.

Outline of this corrigendum

In the next section, we adapt the statement and proof of convergence of our approachability strategy to the class of games of partial monitoring such that the mappings $\overline{m}(\cdot, b)$ are piecewise linear for all $b \in \mathcal{B}$; we also show that the minimization of (external or internal) regrets fall in this case. In a third and last section, we will show how the approachability of a general closed convex set for a general game can be handled.

2 Bi-Piecewise Linear Games — Minimization of Regrets

We first state what the proof of Lemma 10 of the original paper correctly shows.

Lemma 1 *For any game of 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*

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

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

The results of this section will rely on the following assumption.

Assumption 1 *A game is bi-piecewise linear if $\overline{m}(\cdot, b)$ is piecewise linear on $\Delta(\mathcal{I})$ for every $b \in \mathcal{B}$.*

Assumption 1 means that for all $b \in \mathcal{B}$ there exists a decomposition of $\Delta(\mathcal{I})$ into polytopes each on which $\overline{m}(\cdot, b)$ is linear. Since \mathcal{B} is finite, there exists a finite number of such decompositions, and thus there exists a polytopial decomposition that refines all of them. (The latter is generated by the intersection of all considered polytopes as b varies.) By construction, every $\overline{m}(\cdot, b)$ 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 10 then 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

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

where we denoted the convex weight vector $\Theta(\mathbf{p}) \in \Delta(\mathcal{A})$ by $(\Theta_a(\mathbf{p}))_{a \in \mathcal{A}}$. Therefore, on a lifted space, \overline{m} is seen to coincide with a bi-linear mapping.

Definition 2 *We denote by $\overline{\overline{m}}$ the linear extension to $\Delta(\mathcal{A} \times \mathcal{B})$ of the restriction of \overline{m} to $\mathcal{A} \times \mathcal{B}$, so that for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$,*

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

The main argument follows the same lines as Section 4.2.2 of the original paper: we construct an (r, H) -approachability strategy for the original problem based on a strategy for \overline{m} -robust approachability of \mathcal{C} . We do so by considering the following equivalent to Lemma 12 of the original paper. Condition 1 refers to the condition

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

stated in the original paper; we indicated therein that it was necessary and we need to adapt the proof of its sufficiency.

Lemma 3 *Under Condition 1, the closed convex set \mathcal{C} is \overline{m} -robust approachable.*

Proof: We show that Condition (RAC) in Theorem 7 of the original paper is satisfied, that is, that for all $\mathbf{y} \in \Delta(\mathcal{B})$, there exists some $\mathbf{x} \in \Delta(\mathcal{A})$ such that $\overline{m}(\mathbf{x}, \mathbf{y}) \subseteq \mathcal{C}$. With a given such $\mathbf{y} \in \Delta(\mathcal{B})$, we associate the feasible vector of signals $\sigma = \sum_{b \in \mathcal{B}} y_b b$ and let \mathbf{p} be given by Condition 1, so¹ that $\overline{m}(\mathbf{p}, \sigma) \subseteq \mathcal{C}$. By linearity of \overline{m} (for the first equality), by definition of \overline{m} (for the first inclusion), by Lemma 1 (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 now indicate how our algorithm (stated in Figure 1 of the original paper) should be slightly adapted; we only write the steps that need a modification.

[...]

Initialization: [...] 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 $\mathbf{p}_n = (1 - \gamma) \mathbf{x}_n + \gamma \mathbf{u}$;

[...]

5. choose $\theta_{n+1} = \Psi(\theta_1, \Phi(\hat{\sigma}_1), \dots, \theta_n, \Phi(\hat{\sigma}_n))$.
-

Figure 1: The modifications to perform on the proposed strategy.

The proof that this strategy indeed approaches \mathcal{C} at $T^{-1/5}$ rate with overwhelming probability is adapted as follows from the proof of Theorem 13, which can be found in the extended arXiv version of the original paper. The approximation and concentration results stated in Equations (9)–(11) remain unchanged, so that

$$\frac{1}{T} \sum_{t=1}^T r(I_t, J_t) \quad \text{is close to} \quad \frac{1}{N} \sum_{n=1}^N r(\mathbf{x}_n, \hat{\mathbf{q}}_n) = \frac{1}{N} \sum_{n=1}^N \sum_{a \in \mathcal{A}} \theta_{n,a} r(a, \hat{\mathbf{q}}_n).$$

Now, by definition of \overline{m} ,

$$\frac{1}{N} \sum_{n=1}^N \sum_{a \in \mathcal{A}} \theta_{n,a} r(a, \hat{\mathbf{q}}_n) \quad \text{belongs to the set} \quad \frac{1}{N} \sum_{n=1}^N \sum_{a \in \mathcal{A}} \theta_{n,a} \overline{m}(a, \tilde{H}(\hat{\mathbf{q}}_n)).$$

By definition of Φ and by linearity of \overline{m} ,

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

¹Note however that we do not necessarily have that $\Phi(\sigma)$ and \mathbf{y} are equal, as Φ is not a one-to-one mapping.

Equation (12) can be adapted along the same lines as in the original paper, since its proof only relied on a Lipschitzness property and a concentration argument:

$$\frac{1}{N} \sum_{n=1}^N \bar{m} \left(\theta_n, \Phi \left(\tilde{H}(\hat{\mathbf{q}}_n) \right) \right) \quad \text{is close to the set} \quad \frac{1}{N} \sum_{n=1}^N \bar{m} \left(\theta_n, \Phi(\hat{\sigma}_n) \right).$$

Finally, since Ψ is a strategy designed for the \bar{m} -robust approachability of \mathcal{C} ,

$$\frac{1}{N} \sum_{n=1}^N \bar{m} \left(\mathbf{x}_n, \Phi(\hat{\sigma}_n) \right) \quad \text{gets closer and closer to the set } \mathcal{C} \text{ when } n \rightarrow \infty,$$

which concludes this (sketch of) proof.

Minimization of Regrets

It only remains to indicate why the results of Section 5 of the original paper (on minimization of regrets) still hold. For the case of external regret, for instance, we can find a convex set \mathcal{C} and a vector-valued payoff function \underline{r} , that first, satisfy Assumption 1 and second, are such that R_T^{ext} is still upper bounded by (a constant depending on the game only times) the distance of the average payoff vector $(1/T) \sum_{t \leq T} \underline{r}(I_t, J_t)$ to \mathcal{C} .

These are indeed given by

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

for all $(i, j) \in \mathcal{I} \times \mathcal{J}$. That (\underline{r}, H) -approachability of \mathcal{C} entails minimization of the regret R_T^{ext} is straightforward along the same lines as the first part of the proof of Corollary 14 in the original paper.

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

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

is piecewise linear. Since \tilde{H} is linear, the set

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

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

$$\bar{m}_1(\mathbf{p}, \sigma) = \text{co} \left\{ r(\mathbf{p}, \mathbf{q}_{\sigma,1}), \dots, r(\mathbf{p}, \mathbf{q}_{\sigma,M}) \right\} = \left[\begin{array}{c} \min_{k \in \{1, \dots, M\}} r(\mathbf{p}, \mathbf{q}_{\sigma,k}), \quad \max_{k' \in \{1, \dots, M\}} r(\mathbf{p}, \mathbf{q}_{\sigma,k'}) \end{array} \right],$$

where co stands for the convex hull. Since all applications $r(\cdot, \mathbf{q}_{\sigma,k})$ are linear, their minimum and their maximum are piecewise linear functions, thus $\bar{m}_1(\cdot, \sigma)$ is also piecewise linear.

For the internal regret, the bi-piecewise linearity of the game (up to the same slight modification of the payoff function \underline{r} and of the definition of \mathcal{C}) follows from a similar argument.

3 The Case of General Games

Unfortunately, as the example below illustrates, there exist game that are not bi-piecewise linear.

Example 2 Consider the following game.

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

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 $\overline{m}((p, 1 - p), \clubsuit)$ 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 1 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 $\overline{p} = \lambda p + (1 - \lambda)p'$, one necessarily has that

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

the inclusion \subseteq holds by definition of \overline{m} 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.

However, we will show that if Condition 1 holds there exist strategies with a constant per-round complexity to approach polytopes even when the game is not bi-piecewise linear. That is, by considering simpler convex sets \mathcal{C} , no assumption is needed on the pair (r, H) . We will conclude this note by indicating that thanks to a doubling trick, Condition 1 is still seen to be sufficient for approachability in the most general case when no assumption is made neither on (r, H) nor on \mathcal{C} , at the cost however of inefficiency.

Approachability of the Negative Orthant in the Case of General Games

For the sake of simplicity, we start with the case of the negative orthant \mathbb{R}_-^d . Our argument will be based on Lemma 1; 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 into 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

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

where Φ refers to the mapping defined in Lemma 1. The lemma below indicates why \check{m} is an excellent substitute to \overline{m} in the case of the approachability of the orthant \mathbb{R}_-^d .

Lemma 4 *The set-valued mappings \check{m} and \overline{m} are linked by the following two properties: for all $\mathbf{p} \in \Delta(\mathcal{I})$ and $\sigma \in \mathcal{F}$,*

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

The interpretations of these two properties are that 1. \check{m} -robust approaching a set \mathcal{C} is more difficult than \overline{m} -robust approaching it; and 2. that if Condition 1 holds for \overline{m} and \mathbb{R}_-^d , it also holds for \check{m} and \mathbb{R}_-^d .

Proof: For property 1., note that by construction of \check{m} ,

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

Lemma 1 and the linear extension of \tilde{m} then show that

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

As for property 2., it suffices to note that (by Lemma 1 again) the stated assumption exactly means that $\sum_{b \in \mathcal{B}} \Phi_b(\sigma) \overline{m}(\mathbf{p}, b) \subset \mathbb{R}_-^d$. 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 as announced; note that no bi-piecewise linearity assumption is needed on the game.

Theorem 5 *If Condition 1 is satisfied for \bar{m} , 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.*

Proof: We will apply the result of the previous section. By checking the proof of the main theorem (Theorem 13) of the original paper, one can see that the only ingredient needed is a strategy for \bar{m} -robust approaching $\mathcal{C} = \mathbb{R}_-^d$. But by Lemma 4, it is therefore enough to \tilde{m} -robust approach $\mathcal{C} = \mathbb{R}_-^d$. Now, a strategy performing this exists because of the result stated in Theorem 13 and corrected in the previous section, as first, Condition 1 holds for \tilde{m} as well (as indicated by Lemma 4) and second, \tilde{m} is bi-piecewise linear. The latter fact can be seen by showing —similarly to what was done in the section devoted to regret minimization— that each \tilde{m}_k , thus also \tilde{m} , is piecewise linear. ■

Approachability of Polytopes in the Case of General Games

If that 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 an exercise of mere rewriting.

Lemma 6 *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 all strategies for one problem translates to a strategy for the other problem.*

In addition, Condition 1 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 5 indicates that Condition 1 for (r, H) and \mathcal{C} is a sufficient condition for the (r, H) -approachability of \mathcal{C} .

Approachability of General Convex Sets in the Case of General Games

A general closed convex set can always be approximated arbitrarily well by a polytope (where the number of vertices of the latter however increases as the quality of the approximation does). There, via a doubling trick, Condition 1 is also seen to be sufficient to (r, H) -approach any general closed convex set \mathcal{C} . However, the computational complexity of the resulting strategy is much larger: the per-round complexity increases over time (as the numbers of vertices of the approximating polytopes do).