

Lecture 5

Learning in Games: Algorithms

Instructor: Gabriele Farina*

In Lecture 4 we have mentioned how no-external-regret dynamics recover several solution concepts of interest, including Nash equilibria in two-player zero-sum games, normal-form coarse-correlated equilibria in multiplayer general-sum games, and more generally convex-concave saddle point problems. Furthermore, we have seen that no-external-regret dynamics are a fundamental building block for constructing no- Φ -regret dynamics for general sets Φ of deviation functions.

In this lecture, we start exploring how no-external-regret dynamics for players in normal-form games can be constructed. As a reminder, a player with n action has a strategy space that corresponds to the probability simplex of all distributions over those actions,

$$\Delta^m = \{(x_1, \dots, x_m) \in \mathbb{R}_{\geq 0}^m : x_1 + \dots + x_m = 1\}.$$

Furthermore, remember that constructing a regret minimizer for Δ^m means that we need to build a mathematical object that supports two operations:

- `NextStrategy()` has the effect that the regret minimizer will output an element $\mathbf{x}^{(t)} \in \Delta^m$;
- `ObserveUtility($u^{(t)}$)` provides the environment's feedback to the regret minimizer, in the form of a utility function $u^{(t)}$. For today, we will focus on a case of particular relevance for normal-form games: the case of *linear* utilities (as such are the utilities in a normal-form game). In particular, to fix notation, we will let

$$u^{(t)} : \mathbf{x} \mapsto (\mathbf{g}^{(t)})^\top \mathbf{x},$$

where $\mathbf{g}^{(t)} \in \mathbb{R}^n$ is the vector that defines the linear function (called with the letter \mathbf{g} to suggest the idea of it being the “*gradient vector*”). Since the utility in the game cannot be unbounded, we assume that the $\mathbf{g}^{(t)}$ can be arbitrary but *bounded in norm*.

In this notation, the (external) *regret* is defined as the quantity

$$\text{Reg}^{(T)} := \max_{\hat{\mathbf{x}} \in \Delta^m} \left\{ \sum_{t=1}^T \left(\langle \mathbf{g}^{(t)}, \hat{\mathbf{x}} \rangle - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \right) \right\}.$$

Our goal is to make sure that the regret grows sublinearly in T no matter the utility vectors $\mathbf{g}^{(t)}$ chosen by the environment.

Most algorithms known today for the task operate by *prioritizing actions based on how much regret they have incurred*. In particular, a key quantity to define modern algorithms is the vector of cumulated actions regrets

$$\mathbf{r}^{(t)} := \sum_{\tau=1}^t \mathbf{g}^{(\tau)} - \langle \mathbf{g}^{(\tau)}, \mathbf{x}^{(\tau)} \rangle \mathbf{1}.$$

(Note that $\text{Reg}^{(t)} = \max_{i \in \{1, \dots, m\}} \mathbf{r}^{(t)}[i]$.) The natural question is: how to prioritize?

*MIT EECS. ✉ gfarina@mit.edu.

1 An algorithm that does NOT work: Follow-the-Leader

Perhaps the most natural idea would be to always play the action with the highest cumulated regret (breaking ties, say, lexicographically). After all, this is the action we wish the most we had played in the past. Unfortunately, this idea is known not to work. To see that, consider the following sequence of gradient vectors:

$$\mathbf{g}^{(1)} = \begin{pmatrix} 0 \\ 1/2 \end{pmatrix}, \quad \mathbf{g}^{(2)} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \mathbf{g}^{(3)} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \mathbf{g}^{(4)} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \mathbf{g}^{(5)} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \text{etc.}$$

In this case, at time 1, we pick the first action and score a utility of 0. We then compute

$$\mathbf{r}^{(1)} = \begin{pmatrix} 0 \\ 1/2 \end{pmatrix}.$$

and at time $t = 2$ pick the second action since it has the highest regret. This results in a score of 0 and an updated regret vector

$$\mathbf{r}^{(2)} = \begin{pmatrix} 1 \\ 1/2 \end{pmatrix}$$

So, at time $t = 3$, we will pick the first action, resulting again in a score of 0 and an updated regret

$$\mathbf{r}^{(3)} = \begin{pmatrix} 1 \\ 3/2 \end{pmatrix},$$

and so on... Overall, it's easy to see that in all this jumping around, the regrets grow linearly.

2 The Regret Matching (RM) algorithm

The regret matching algorithm picks probabilities proportional to the “ReLU” of the cumulated regrets, that is,

$$\mathbf{x}^{(t+1)} \propto [\mathbf{r}^{(t)}]^+ \tag{1}$$

whenever $[\mathbf{r}^{(t)}] \neq \mathbf{0}$, and an arbitrary point otherwise.

The algorithm is presented in pseudocode in Algorithm 1.

Theorem 2.1. The regret matching algorithm (Algorithm 1) is an external regret minimizer, and satisfies the regret bound $R^T \leq \Omega\sqrt{T}$, where Ω is the maximum norm $\|\mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}\|_2$ up to time T .

Proof. We start by observing that, at all times t ,

$$\left(\mathbf{g}^{(t+1)} - \langle \mathbf{g}^{(t+1)}, \mathbf{x}^{(t+1)} \rangle \mathbf{1} \right)^\top \mathbf{x}^{(t+1)} = 0$$

(this is always true, not just for regret matching). Plugging in the definition (1) of how $\mathbf{x}^{(t+1)}$ is constructed, we therefore conclude that

$$\left(\mathbf{g}^{(t+1)} - \langle \mathbf{g}^{(t+1)}, \mathbf{x}^{(t+1)} \rangle \mathbf{1} \right)^\top [\mathbf{r}^{(t)}]^+ = 0. \tag{2}$$

(Note that the above equation holds trivially when $[\mathbf{r}^{(t)}]^+$ and therefore $\mathbf{x}^{(t+1)}$ is picked arbitrarily.)

Now, we use the following inequality:

$$\|[\mathbf{a} + \mathbf{b}]^+\|_2^2 \leq \|[\mathbf{a}]^+ + \mathbf{b}\|_2^2,$$

applied to $\mathbf{a} = \mathbf{r}^{(t)}$ and $\mathbf{b} = \mathbf{g}^{(t+1)} - \langle \mathbf{g}^{(t+1)}, \mathbf{x}^{(t+1)} \rangle \mathbf{1}$. In particular, since by definition $\mathbf{a} + \mathbf{b} = \mathbf{r}^{(t+1)}$, we have

$$\begin{aligned} \|[\mathbf{r}^{(t+1)}]^+\|_2^2 &\leq \|[\mathbf{r}^{(t)}]^+ + (\mathbf{g}^{(t+1)} - \langle \mathbf{g}^{(t+1)}, \mathbf{x}^{(t+1)} \rangle \mathbf{1})\|_2^2 \\ &= \|[\mathbf{r}^{(t)}]^+\|_2^2 + \|\mathbf{g}^{(t+1)} - \langle \mathbf{g}^{(t+1)}, \mathbf{x}^{(t+1)} \rangle \mathbf{1}\|_2^2 + 2\left(\mathbf{g}^{(t+1)} - \langle \mathbf{g}^{(t+1)}, \mathbf{x}^{(t+1)} \rangle \mathbf{1}\right)^\top [\mathbf{r}^{(t)}]^+ \\ &= \|[\mathbf{r}^{(t)}]^+\|_2^2 + \|\mathbf{g}^{(t+1)} - \langle \mathbf{g}^{(t+1)}, \mathbf{x}^{(t+1)} \rangle \mathbf{1}\|_2^2 && \text{(from (2))} \\ &\leq \|[\mathbf{r}^{(t)}]^+\|_2^2 + \Omega^2. \end{aligned}$$

Hence, by induction we have

$$\|[\mathbf{r}^{(T)}]^+\|_2^2 \leq T\Omega^2 \implies \text{Reg}^{(T)} = \max_{i \in \{1, \dots, m\}} \mathbf{r}^{(T)}[i] \leq \max_{i \in \{1, \dots, m\}} [\mathbf{r}^{(T)}[i]]^+ \leq \|[\mathbf{r}^{(T)}]^+\|_2 \leq \Omega\sqrt{T}.$$

The proof is then complete. \square

Algorithm 1: Regret matching

```

1  $\mathbf{r}^0 \leftarrow \mathbf{0} \in \mathbb{R}^n, \mathbf{x}^0 \leftarrow \mathbf{1}/m \in \Delta^m$ 
2 function NextStrategy()
3   if  $[\mathbf{r}^{(t-1)}]^+ \neq \mathbf{0}$  return  $\mathbf{x}^{(t)} \leftarrow \frac{[\mathbf{r}^{(t-1)}]^+}{\|[\mathbf{r}^{(t-1)}]^+\|_1}$ 
4   else return  $\mathbf{x}^{(t)} \leftarrow$  any point in  $\Delta^m$ 
5 function ObserveUtility( $\mathbf{g}^t$ )
6    $\mathbf{r}^{(t)} \leftarrow \mathbf{r}^{(t-1)} + \mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}$ 

```

Algorithm 2: Regret matching⁺

```

1  $\mathbf{z}^0 \leftarrow \mathbf{0} \in \mathbb{R}^n, \mathbf{x}^0 \leftarrow \mathbf{1}/m \in \Delta^m$ 
2 function NextStrategy()
3   if  $[\mathbf{r}^{(t-1)}]^+ \neq \mathbf{0}$  return  $\mathbf{x}^{(t)} \leftarrow \frac{[\mathbf{r}^{(t-1)}]^+}{\|[\mathbf{r}^{(t-1)}]^+\|_1}$ 
4   else return  $\mathbf{x}^{(t)} \leftarrow$  any point in  $\Delta^m$ 
5 function ObserveUtility( $\mathbf{g}^t$ )
6    $\mathbf{z}^{(t)} \leftarrow [\mathbf{z}^{(t-1)} + \mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}]^+$ 

```

As of today, regret matching and its variants are still often some of the most practical algorithms for learning in games.

Remark 2.1. One very appealing property of the regret matching algorithm is its *lack of hyperparameters*. It just works “out of the box”.

2.1 The regret matching⁺ (RM⁺) algorithm

The regret matching⁺ algorithm was introduced by Tammelin [2014], Tammelin et al. [2015], and is given in Algorithm 2. It differs from RM only on the last line, where a further thresholding is added. That small change has the effect that actions with negative cumulated regret (that is, “bad” actions) are treated as actions with 0 regret. Hence, intuitively, if a bad action were to become good over time, it would take less time for RM⁺ to notice and act on that change. Because of that, regret matching⁺ has stronger practical performance and is often preferred over regret matching in the game solving literature.

With a simple modification to the analysis of RM, the same bound as RM can be proven.

Theorem 2.2. The RM^+ algorithm (Algorithm 2) is an external regret minimizer, and satisfies the regret bound $R^T \leq \Omega\sqrt{T}$, where Ω is the maximum norm $\|\mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}\|_2$ up to time T .

3 Follow-the-Regularized-Leader (FTRL)

Another way of obtaining no-regret algorithms is by considering a regularized (i.e., smoothed) version of the follow-the-leader algorithm discussed above. The idea is that, instead of playing by always putting 100% of the probability mass on the action with highest cumulated regret, we look for the distribution that maximizes the expected cumulated regret, *minus* some regularization term that prevents us from putting all the mass on a single action.

More specifically, let $\varphi : \Delta^m \rightarrow \mathbb{R}$ be a 1-strongly convex¹ function (with respect to some norm $\|\cdot\|$, for example, the Euclidean norm $\|\cdot\|_2$) with bounded range, and consider the choice of strategy

$$\mathbf{x}^{(t)} := \arg \max_{\hat{\mathbf{x}} \in \Delta^m} \left\{ (\mathbf{r}^{(t-1)})^\top \hat{\mathbf{x}} - \frac{1}{\eta} \varphi(\hat{\mathbf{x}}) \right\}.$$

(so, in particular $\mathbf{x}^{(1)} := \arg \min_{\hat{\mathbf{x}} \in \Delta^m} \varphi(\hat{\mathbf{x}})$ since the regrets of all actions are 0 at the beginning). This algorithm is called the *follow-the-regularized-leader (FTRL)* algorithm. The following two properties are known about FTRL:

- Regularization *limits the amount of variation between consecutive strategies*. This makes intuitive sense: if $\eta \rightarrow 0$, for example, $\mathbf{x}^{(t)}$ is constant (and equal to $\arg \min_{\hat{\mathbf{x}} \in \Delta^m} \varphi(\hat{\mathbf{x}})$). When $\eta = \infty$, we recover the follow-the-leader algorithm, where strategies can jump arbitrarily. For intermediate η , as you might expect, the amount of variation between strategies at consecutive times is bounded above by a quantity proportional to η :

$$\|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\| \leq \eta \|\mathbf{g}^{(t)}\|_*, \quad (3)$$

where $\|\cdot\|_*$ is the *dual* norm of $\|\cdot\|$ (if $\|\cdot\|$ is the Euclidean norm, its dual is also the Euclidean norm; if $\|\cdot\|$ is the ℓ_1 norm, then its dual is the ℓ_∞ norm).

- A useful regret bound of FTRL is given by

$$\text{Reg}^{(T)} \leq \max_{\hat{\mathbf{x}} \in \Delta^m} \frac{\varphi(\hat{\mathbf{x}}) - \varphi(\mathbf{x}^{(1)})}{\eta} + \sum_{t=1}^T (\mathbf{g}^{(t)})^\top (\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}). \quad (4)$$

Using the Cauchy-Schwarz inequality in (4) and plugging in (3), we obtain the following

Theorem 3.1. The regret cumulated by the FTRL algorithm is upper bounded by

$$\text{Reg}^{(T)} \leq \max_{\hat{\mathbf{x}} \in \Delta^m} \frac{\varphi(\hat{\mathbf{x}}) - \varphi(\mathbf{x}^{(1)})}{\eta} + \eta \sum_{t=1}^T \|\mathbf{g}^{(t)}\|_*^2.$$

Since the norm of the gradient vectors is bounded, it follows immediately that by picking learning rate $\eta = 1/\sqrt{T}$, we obtain that $\text{Reg}^{(T)}$ is bounded as roughly \sqrt{T} (a sublinear function!) at all times T .

One of the appeals of FTRL is that it is a very general method, and in fact it generalized well beyond probability simplexes.

We will now see a particular instantiation of FTRL that is so important and well-studied that it deserves its own name: the *Multiplicative Weights Update* method (MWU).

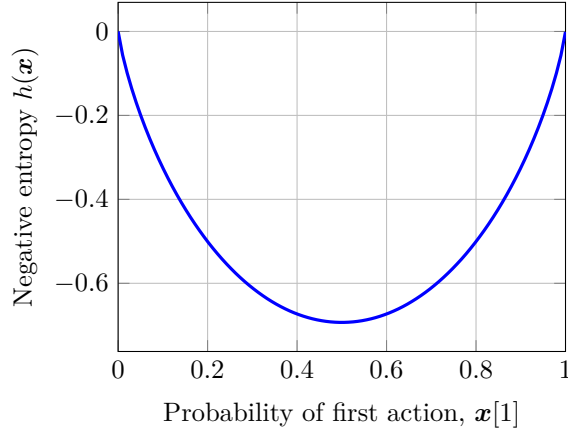
¹A differentiable function f is 1-strongly convex if for all \mathbf{x}, \mathbf{y} in its domain one has $(\nabla f(\mathbf{x}) - \nabla f(\mathbf{y}))^\top (\mathbf{x} - \mathbf{y}) \geq \|\mathbf{x} - \mathbf{y}\|^2$.

3.1 The Multiplicative Weights Update Algorithm

A natural choice of convex regularizer for probability simplexes is the *negative entropy* function

$$h(\mathbf{x}) := \sum_{i \in \{1, \dots, m\}} \mathbf{x}[i] \log \mathbf{x}[i].$$

Example 3.1. For example, this is what negative entropy looks like in the case of $m = 2$ actions:



The negative entropy function has the following properties:

- It is 1-strongly convex with respect to the ℓ_1 norm $\|\cdot\|_1$.
- The maximum of the function is 0, which is attained at any deterministic strategy.
- Its minimum (that is, maximum entropy) is attained at the uniformly random strategy $\bar{\mathbf{x}} = (1/m, \dots, 1/m)$, at which the function has value $h(\bar{\mathbf{x}}) = m \cdot 1/m \log 1/m = -\log m$.

In this case, then, the MWU algorithm reduces to the following:

$$\begin{aligned} \mathbf{x}^{(1)} &:= (1/m, \dots, 1/m), \\ \mathbf{x}^{(t)} &:= \arg \max_{\hat{\mathbf{x}} \in \Delta^m} \left\{ (\mathbf{r}^{(t-1)})^\top \hat{\mathbf{x}} - \frac{1}{\eta} h(\hat{\mathbf{x}}) \right\} = \arg \max_{\hat{\mathbf{x}} \in \Delta^m} \left\{ (\mathbf{r}^{(t-1)})^\top \hat{\mathbf{x}} - \frac{1}{\eta} \sum_{i=1}^m \mathbf{x}[i] \log \mathbf{x}[i] \right\}. \end{aligned} \quad (5)$$

As it turns out via a simple application of the Lagrange multiplier theorem, the concave optimization problem in (5) has a closed-form solution:

Proposition 3.1. The concave optimization problem (5) has the closed-form solution

$$\mathbf{x}^{(t)} = \text{softmax}\{\eta \mathbf{r}^{(t)}\},$$

that is, for every action $i \in \{1, \dots, m\}$,

$$\mathbf{x}^{(t)}[i] = \frac{e^{\eta \mathbf{r}^{(t)}[i]}}{\sum_{j=1}^m e^{\eta \mathbf{r}^{(t)}[j]}}.$$

Algorithm 3: Multiplicative Weights Update

```
1  $\mathbf{r}^0 \leftarrow \mathbf{0} \in \mathbb{R}^n, \mathbf{x}^0 \leftarrow \mathbf{1}/m \in \Delta^m$   
2 function NextStrategy()  
3 |   return  $\mathbf{x}^{(t)} = \text{softmax}(\eta \mathbf{r}^{(t-1)})$   
4 function ObserveUtility( $\mathbf{g}^t$ )  
5 |    $\mathbf{r}^{(t)} \leftarrow \mathbf{r}^{(t-1)} + \mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}$ 
```

This gives rise to Algorithm 3, which has a similar flavor as RM, but uses a different prioritization of actions (in particular, it uses *softmax instead of ReLU*). From the general analysis of FTRL (Theorem 3.1), remembering that:

- the maximum of h is 0,
- the minimum (attained in $\mathbf{x}^{(1)}$) is $-\log m$,
- negative entropy is 1-strongly convex with respect to $\|\cdot\|_1$ and that the dual norm to $\|\cdot\|_1$ is $\|\cdot\|_\infty$,

we obtain the following.

Corollary 3.1. The regret cumulated by the MWU algorithm is upper bounded by

$$\text{Reg}^{(T)} \leq \frac{\log m}{\eta} + \eta \sum_{t=1}^T \|\mathbf{g}^{(t)}\|_\infty^2.$$

References

- Oskari Tammelin. Solving large imperfect information games using CFR+, 2014.
- Oskari Tammelin, Neil Burch, Michael Johanson, and Michael Bowling. Solving heads-up limit Texas hold'em. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150, 2000.
- David Blackwell. An analog of the minmax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6: 1–8, 1956.

Appendix

The regret matching (RM) algorithm has a connection to an ancient (and very elegant!) construction, called *Blackwell approachability*. Blackwell approachability is a precursor of the theory of regret minimization, and played a fundamental role in the historical development of several efficient online optimization methods. In particular, as we will show in a minute, the problem of minimizing regret on a simplex can be rewritten as a Blackwell approachability game. The solution of the Blackwell approachability game will then recover exactly RM.

A Blackwell approachability game

Blackwell approachability generalizes the problem of playing a repeated two-player game to games whose utilities are vectors instead of scalars.

Definition A.1. A Blackwell approachability game is a tuple $(\mathcal{X}, \mathcal{Y}, \mathbf{f}, S)$, where \mathcal{X}, \mathcal{Y} are closed convex sets, $\mathbf{f} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ is a biaffine function, and $S \subseteq \mathbb{R}^d$ is a closed and convex *target set*. A Blackwell approachability game represents a vector-valued repeated game between two players. At each time t , the two players interact in this order:

- first, Player 1 selects an action $\mathbf{x}^{(t)} \in \mathcal{X}$;
- then, Player 2 selects an action $\mathbf{y}^{(t)} \in \mathcal{Y}$, which can depend adversarially on all the $\mathbf{x}^{(t)}$ output so far;
- finally, Player 1 incurs the vector-valued payoff $\mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \in \mathbb{R}^d$, where \mathbf{f} is a biaffine function.

Player 1's objective is to guarantee that the average payoff converges to the target set S . Formally, given target set $S \subseteq \mathbb{R}^d$, Player 1's goal is to pick actions $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots \in \mathcal{X}$ such that no matter the actions $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots \in \mathcal{Y}$ played by Player 2,

$$\min_{\hat{\mathbf{s}} \in S} \left\| \hat{\mathbf{s}} - \frac{1}{T} \sum_{t=1}^T \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \right\|_2 \rightarrow 0 \quad \text{as } T \rightarrow \infty. \quad (6)$$

B Regret minimization on the simplex via Blackwell approachability

Hart and Mas-Colell [2000] noted that the construction of a regret minimizer for a simplex domain Δ^n can be reduced to constructing an algorithm for a particular Blackwell approachability game $\Gamma := (\Delta^n, \mathbb{R}^n, \mathbf{f}, \mathbb{R}_{\leq 0}^n)$ which we now describe. For all $i \in \{1, \dots, n\}$, the i -th component of the vector-valued payoff function \mathbf{f} measures the change in regret incurred at time t , compared to always playing the i -th vertex \mathbf{e}_i of the simplex. Formally, $\mathbf{f} : \Delta^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is defined as

$$\mathbf{f}(\mathbf{x}^{(t)}, \mathbf{g}^{(t)}) := \mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}, \quad (7)$$

where $\mathbf{1}$ is the n -dimensional vector whose components are all 1. (Note the connection with $\mathbf{r}^{(t)}$ seen in Lecture 5).

The following lemma establishes an important link between Blackwell approachability on Γ and external regret minimization on the simplex Δ^n .

Lemma B.1. The regret $\text{Reg}^{(T)} = \max_{\hat{\mathbf{x}} \in \Delta^n} \frac{1}{T} \sum_{t=1}^T \langle \mathbf{g}^{(t)}, \hat{\mathbf{x}} - \mathbf{x}^{(t)} \rangle$ cumulated up to any time T by any sequence of decisions $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)} \in \Delta^n$ is related to the distance of the average Blackwell payoff from the target cone $\mathbb{R}_{\leq 0}^n$ as

$$\frac{\text{Reg}^{(T)}}{T} \leq \min_{\hat{\mathbf{s}} \in \mathbb{R}_{\leq 0}^n} \left\| \hat{\mathbf{s}} - \frac{1}{T} \sum_{t=1}^T \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{g}^{(t)}) \right\|_2. \quad (8)$$

So, a strategy for the Blackwell approachability game Γ is a regret-minimizing strategy for the simplex domain Δ^n .

Proof. For any $\hat{\mathbf{x}} \in \Delta^n$, the regret cumulated compared to always playing $\hat{\mathbf{x}}$ satisfies

$$\begin{aligned} \frac{1}{T} \text{Reg}^{(T)}(\hat{\mathbf{x}}) &:= \frac{1}{T} \sum_{t=1}^T \left(\langle \mathbf{g}^{(t)}, \hat{\mathbf{x}} \rangle - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \right) = \frac{1}{T} \sum_{t=1}^T \left(\langle \mathbf{g}^{(t)}, \hat{\mathbf{x}} \rangle - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \langle \mathbf{1}, \hat{\mathbf{x}} \rangle \right) \\ &= \left\langle \frac{1}{T} \sum_{t=1}^T \mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}, \hat{\mathbf{x}} \right\rangle = \left\langle \frac{1}{T} \sum_{t=1}^T \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{g}^{(t)}), \hat{\mathbf{x}} \right\rangle \\ &= \min_{\hat{\mathbf{s}} \in \mathbb{R}_{\leq 0}^n} \left\langle -\hat{\mathbf{s}} + \frac{1}{T} \sum_{t=1}^T \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{g}^{(t)}), \hat{\mathbf{x}} \right\rangle, \end{aligned} \quad (9)$$

where we used the fact that $\hat{\mathbf{x}} \in \Delta^n$ in the second equality, and that $\min_{\hat{\mathbf{s}} \in \mathbb{R}_{\leq 0}^n} \langle -\hat{\mathbf{s}}, \hat{\mathbf{x}} \rangle = 0$ since $\hat{\mathbf{x}} \geq \mathbf{0}$. Applying the Cauchy-Schwarz inequality to the right-hand side of (9), we obtain

$$\frac{1}{T} \text{Reg}^{(T)}(\hat{\mathbf{x}}) \leq \min_{\hat{\mathbf{s}} \in \mathbb{R}_{\leq 0}^n} \left\| -\hat{\mathbf{s}} + \frac{1}{T} \sum_{t=1}^T \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{g}^{(t)}) \right\|_2 \|\hat{\mathbf{x}}\|_2.$$

So, using the fact that $\|\hat{\mathbf{x}}\|_2 \leq 1$ for any $\hat{\mathbf{x}} \in \Delta^n$,

$$\frac{1}{T} \text{Reg}^{(T)}(\hat{\mathbf{x}}) \leq \min_{\hat{\mathbf{s}} \in \mathbb{R}_{\leq 0}^n} \left\| -\hat{\mathbf{s}} + \frac{1}{T} \sum_{t=1}^T \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{g}^{(t)}) \right\|_2.$$

Taking a max over $\hat{\mathbf{x}} \in \Delta^n$ yields the statement. \square

C Solving Blackwell games: Blackwell's algorithm

A central concept in the theory of Blackwell approachability is the following.

Definition C.1 (Forceable halfspace). Let $(\mathcal{X}, \mathcal{Y}, \mathbf{f}, S)$ be a Blackwell approachability game and let $\mathcal{H} \subseteq \mathbb{R}^d$ be a halfspace, that is, a set of the form $\mathcal{H} = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{a}^\top \mathbf{x} \leq b\}$ for some $\mathbf{a} \in \mathbb{R}^d, b \in \mathbb{R}$. The halfspace \mathcal{H} is said to be *forceable* if there exists a strategy of Player 1 that guarantees that the payoff is in \mathcal{H} no matter the actions played by Player 2, that is, if there exists $\mathbf{x}^* \in \mathcal{X}$ such that

$$\mathbf{f}(\mathbf{x}^*, \mathbf{y}) \in \mathcal{H} \quad \forall \mathbf{y} \in \mathcal{Y}.$$

When that is the case, we call action \mathbf{x}^* a *forcing action* for \mathcal{H} .

Blackwell's *approachability theorem* [Blackwell, 1956] states the following.

Theorem C.1 (Blackwell's theorem). Goal (6) can be attained if and only if every halfspace $H \supseteq S$ is forceable.

We constructively prove the direction that shows how forceability translates into a sequence of strategies that guarantees that goal (6) is attained. Let $(\mathcal{X}, \mathcal{Y}, \mathbf{f}, S)$ be the Blackwell game. The method is pretty simple: at each time step $t = 1, 2, \dots$ operate the following:

1. Compute the average payoff received so far, that is, $\phi^{(t)} = \frac{1}{t} \sum_{\tau=1}^{t-1} \mathbf{f}(\mathbf{x}^{(\tau)}, \mathbf{y}^{(\tau)})$.
2. Compute the Euclidean projection $\psi^{(t)}$ of $\phi^{(t)}$ onto the target set S .

3. If $\phi^{(t)} \in S$ (that is, goal (6) has already been met), pick and play any $\mathbf{x}^{(t)} \in \mathcal{X}$, observe the opponent's action \mathbf{y}^t , and return.
4. Else, consider the halfspace $\mathcal{H}^{(t)}$ tangent to S at the projection point $\psi^{(t)}$, that contains S . In symbols,
$$\mathcal{H}^{(t)} := \{z \in \mathbb{R}^d : (\phi^{(t)} - \psi^{(t)})^\top z \leq (\phi^{(t)} - \psi^{(t)})^\top \psi^{(t)}\}.$$
5. By hypothesis, $\mathcal{H}^{(t)}$ is forceable. Pick $\mathbf{x}^{(t)}$ to be a forcing action for $\mathcal{H}^{(t)}$, observe the opponent's action $\mathbf{y}^{(t)}$, and return.

The above method is summarized in Figure 1.




Figure 1: Construction of the approachability strategy described in Appendix C.

Let's see how the average payoff $\phi^{(t)}$ changes when we play as described above. Clearly,

$$\phi^{(t+1)} = \frac{1}{t} \sum_{\tau=1}^{(t)} \mathbf{f}(\mathbf{x}^{(\tau)}, \mathbf{y}^{(\tau)}) = \frac{t-1}{t} \phi^{(t)} + \frac{1}{t} \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}).$$

Hence, denoting with $\rho^{(t+1)}$ the squared Euclidean distance between $\phi^{(t+1)}$ and the target set, that is,

$$\rho^{(t)} := \min_{\hat{s} \in S} \|\hat{s} - \phi^{(t)}\|_2^2,$$

we have

$$\begin{aligned} \rho^{(t+1)} &\leq \|\psi^{(t)} - \phi^{(t+1)}\|_2^2 = \left\| \psi^{(t)} - \frac{t-1}{t} \phi^{(t)} - \frac{1}{t} \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \right\|_2^2 \\ &= \left\| \frac{t-1}{t} (\psi^{(t)} - \phi^{(t)}) + \frac{1}{t} (\psi^{(t)} - \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})) \right\|_2^2 \\ &= \frac{(t-1)^2}{t^2} \rho^{(t)} + \frac{1}{t^2} \|\psi^{(t)} - \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})\|_2^2 + \frac{2(t-1)}{t^2} \langle \psi^{(t)} - \phi^{(t)}, \psi^{(t)} - \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \rangle. \end{aligned} \quad (10)$$

The proof so far does not use any particular assumption about how $\mathbf{x}^{(t)}$ is picked. Here is where that enters the picture. If $\phi^{(t)} \in S$, then $\psi^{(t)} = \phi^{(t)}$ and therefore the last inner product is equal to 0. Otherwise, we have that $\psi^{(t)} - \phi^{(t)} \neq 0$. In that case, $\mathbf{x}^{(t)}$ is constructed by forcing the halfspace $\mathcal{H}^{(t)}$, and therefore, no matter how $\mathbf{y}^{(t)}$ is picked by the opponent we have

$$\begin{aligned} \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \in \mathcal{H}^{(t)} &\iff (\phi^{(t)} - \psi^{(t)})^\top \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \geq (\phi^{(t)} - \psi^{(t)})^\top \psi^{(t)} \\ &\iff \langle \psi^{(t)} - \phi^{(t)}, \psi^{(t)} - \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \rangle \leq 0. \end{aligned}$$

Plugging in the last inequality into (10) and bounding $\|\boldsymbol{\psi}^{(t)} - \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})\|_2^2 \leq \Omega^2$ where Ω^2 is a diameter parameter of the game (which only depends on \mathbf{f} and S), we obtain

$$\rho^{(t+1)} \leq \frac{(t-1)^2}{t^2} \rho^{(t)} + \frac{\Omega^2}{t^2} \implies t^2 \rho^{(t+1)} - (t-1)^2 \rho^{(t)} \leq \Omega^2 \quad \forall t = 1, 2, \dots$$

Summing the inequality above for $t = 0, \dots, T-1$ and removing the telescoping terms, we obtain

$$T^2 \rho^{(T+1)} \leq T \Omega^2 \implies \rho^{(T+1)} \leq \frac{\Omega^2}{T} \implies \min_{\hat{\mathbf{s}} \in S} \left\| \hat{\mathbf{s}} - \frac{1}{T} \sum_{t=1}^T \mathbf{f}(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \right\|_2 \leq \frac{\Omega}{\sqrt{T}}, \quad (11)$$

which implies that the average payoff in the Blackwell game converges to S at a rate of $O(1/\sqrt{T})$.

D The regret matching (RM) algorithm as an instance of Blackwell's algorithm

First, recall from Appendix B that the external regret minimization on the simplex can be solved via the Blackwell game $\Gamma := (\Delta^n, \mathbb{R}^n, \mathbf{f}, \mathbb{R}_{\leq 0}^n)$ where $\mathbf{f} : \Delta^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is defined as

$$\mathbf{f}(\mathbf{x}^{(t)}, \mathbf{g}^{(t)}) = \mathbf{g}^{(t)} - \langle \mathbf{g}^{(t)}, \mathbf{x}^{(t)} \rangle \mathbf{1}, \quad (12)$$

where $\mathbf{1}$ is the n -dimensional vector whose components are all 1. We will solve this Blackwell approachability game using the strategy explained in Appendix C.

Computation of $\boldsymbol{\psi}^t$ (Step 2). Let's start from looking at how to compute the projection $\boldsymbol{\psi}^t$ of $\boldsymbol{\phi}^t$ onto $S = \mathbb{R}_{\leq 0}^n$. Projection onto the nonpositive orthant amounts to a component-wise minimum with 0, that is, $\boldsymbol{\psi}^t = [\boldsymbol{\phi}^t]^-$. Hence,

$$\boldsymbol{\phi}^{(t)} - \boldsymbol{\psi}^{(t)} = [\boldsymbol{\phi}^{(t)}]^+ \implies (\boldsymbol{\phi}^{(t)} - \boldsymbol{\psi}^{(t)})^\top \boldsymbol{\psi}^{(t)} = 0.$$

Halfspace to be forced (Step 4). Following on with Blackwell's algorithm, when $[\boldsymbol{\phi}^{(t)}]^+ \neq \mathbf{0}$, the halfspace to be forced at each time t is

$$\mathcal{H}^{(t)} := \{\mathbf{z} \in \mathbb{R}^n : \langle [\boldsymbol{\phi}^{(t)}]^+, \mathbf{z} \rangle \leq 0\}.$$

Forcing action for $\mathcal{H}^{(t)}$ (Step 5). We now show that a forcing action for $\mathcal{H}^{(t)}$ indeed exists. Remember that by definition, that is an action $\mathbf{x}^* \in \Delta^n$ such that no matter the $\mathbf{g} \in \mathbb{R}^n$, $\mathbf{f}(\mathbf{x}^*, \mathbf{g}) \in \mathcal{H}^{(t)}$. Expanding the definition of $\mathcal{H}^{(t)}$ and \mathbf{f} , we are looking for a $\mathbf{x}^* \in \Delta^n$ such that

$$\begin{aligned} \langle [\boldsymbol{\phi}^{(t)}], \mathbf{g} - \langle \mathbf{g}, \mathbf{x}^* \rangle \mathbf{1} \rangle \leq 0 \quad \forall \mathbf{g} \in \mathbb{R}^n &\iff \langle [\boldsymbol{\phi}^{(t)}], \mathbf{g} \rangle - \langle \mathbf{g}, \mathbf{x}^* \rangle \langle [\boldsymbol{\phi}^{(t)}]^+, \mathbf{1} \rangle \leq 0 & \forall \mathbf{g} \in \mathbb{R}^n \\ &\iff \langle [\boldsymbol{\phi}^{(t)}], \mathbf{g} \rangle - \langle \mathbf{g}, \mathbf{x}^* \rangle \|[\boldsymbol{\phi}^{(t)}]^+\|_1 \leq 0 & \forall \mathbf{g} \in \mathbb{R}^n \\ &\iff \left\langle \mathbf{g}, \frac{[\boldsymbol{\phi}^{(t)}]}{\|[\boldsymbol{\phi}^{(t)}]^+\|_1} \right\rangle - \langle \mathbf{g}, \mathbf{x}^* \rangle \leq 0 & \forall \mathbf{g} \in \mathbb{R}^n \\ &\iff \left\langle \mathbf{g}, \frac{[\boldsymbol{\phi}^{(t)}]}{\|[\boldsymbol{\phi}^{(t)}]^+\|_1} - \mathbf{x}^* \right\rangle \leq 0 & \forall \mathbf{g} \in \mathbb{R}^n. \end{aligned}$$

Note that we are lucky: $[\boldsymbol{\phi}^{(t)}]^+ / \|[\boldsymbol{\phi}^{(t)}]^+\|_1$ is a nonnegative vector whose entries sum to 1. So, the above inequality can be satisfied with equality for the choice

$$\mathbf{x}^* = \frac{[\boldsymbol{\phi}^{(t)}]^+}{\|[\boldsymbol{\phi}^{(t)}]^+\|_1} \in \Delta^n.$$

In other words, we have that Blackwell's algorithm in this case picks

$$\mathbf{x}^{(t+1)} = \frac{[\boldsymbol{\phi}^{(t)}]^+}{\|[\boldsymbol{\phi}^{(t)}]^+\|_1} \in \Delta^n \iff \mathbf{x}^{(t+1)} \propto [\boldsymbol{\phi}^{(t)}]^+ \propto [\mathbf{r}^{(t)}]^+, \text{ where } \mathbf{r}^{(t)} := \sum_{\tau=1}^t \mathbf{g}^{(\tau)} - \langle \mathbf{g}^{(\tau)}, \mathbf{x}^{(\tau)} \rangle \mathbf{1}.$$

This is exactly the regret matching algorithm seen in Lecture 5.