Robust approachability

with applications to regret minimization in games with partial monitoring

Gilles Stoltz

CNRS — École normale supérieure — INRIA, project-team CLASSIC

& HEC Paris

Joint work with Vianney Perchet (ENS Cachan)
and Shie Mannor (Technion)

Blackwell's approachability

For games with full or bandit monitoring

A vector-valued base game

- Finite action sets $\mathcal{A}$ and $\mathcal{B}$
- Sets of distributions over these action sets $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$
- Payoff function $\boldsymbol{r} : \mathcal{A} \times \mathcal{B} \to \mathbb{R}^d$, linearly extended to $\Delta(\mathcal{A}) \times \Delta(\mathcal{B})$
- Target (often closed convex) set $\mathcal{C} \subset \mathbb{R}^d$

Which is repeated

At each round $t \geqslant 1$,

- The decision maker chooses $\boldsymbol{x}_t \in \Delta(\mathcal{A})$ and possibly draws $I_t \sim \boldsymbol{x}_t$ at random
- Nature chooses $\boldsymbol{y}_t \in \Delta(\mathcal{B})$ and possibly draws $J_t \sim \boldsymbol{y}_t$ at random
- The payoffs $\boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t)$ or $\boldsymbol{r}(I_t, J_t)$ are obtained and observed

Summary

In the non-randomized version, players

- sequentially choose $\boldsymbol{x}_t \in \Delta(\mathcal{A})$ and $\boldsymbol{y}_t \in \Delta(\mathcal{B})$
- obtain the payoff $\boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t) \in \mathbb{R}^d$

We will be interested in the average payoff

$$\overline{\boldsymbol{r}}_T = \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t)$$

Note that by concentration inequalities, it has the same behavior as its randomized counterpart

$$\widetilde{\boldsymbol{r}}_T = \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{r}(I_t, J_t)$$

where $I_t \sim \boldsymbol{x}_t$ and $J_t \sim \boldsymbol{y}_t$.

Summary

In the non-randomized version, players

- sequentially choose $\boldsymbol{x}_t \in \Delta(\mathcal{A})$ and $\boldsymbol{y}_t \in \Delta(\mathcal{B})$
- obtain the payoff $\boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t) \in \mathbb{R}^d$

We will be interested in the average payoff

$$\overline{\boldsymbol{r}}_T = \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t)$$

Goals: Approachability of $\mathcal{C}$

Given the closed convex set $\mathcal{C}$, the decision maker aims at approaching $\mathcal{C}$, i.e., at ensuring that

$$\overline{\boldsymbol{r}}_T \longrightarrow \mathcal{C} \,,$$

while Nature's goal is to prevent this convergence.

Comments

Online and adversarial learning: Data (the $\boldsymbol{y}_t$)

- come sequentially;
- cannot be modeled as the realization of a stochastic process.

Interpretation

Some limit multi-criteria optimization is modeled:

- the vector-valued payoff function $\boldsymbol{r}$ gives the values of several quantities to be controlled or optimized;
- the target convex set $\mathcal{C}$ is the set of acceptable vectors of such values, it indicates when all given criteria are met.

Definition

A set $\mathcal{C} \subset \mathbb{R}^d$ is $\boldsymbol{r}$–approachable if there exists a strategy for the decision maker such that for all strategies of Nature,

$$d(\overline{\boldsymbol{r}}_T, \mathcal{C}) = \inf_{c \in \mathcal{C}} \left\| c - \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t) \right\|_2 \longrightarrow 0$$

It turns out that approachability in the repeated game is equivalent to approachability in a one-shot game in which Nature would play first.

Theorem (Blackwell '56)

A closed convex set $\mathcal{C}$ is approachable if and only

$$\forall \boldsymbol{y} \in \Delta(\mathcal{B}), \quad \exists \boldsymbol{x} \in \Delta(\mathcal{A}), \qquad \boldsymbol{r}(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{C}\,.$$

Theorem (Blackwell '56)

A closed convex set $\mathcal{C}$ is approachable if and only

$$\forall\, \boldsymbol{y} \in \Delta(\mathcal{B}), \quad \exists\, \boldsymbol{x} \in \Delta(\mathcal{A}), \qquad \boldsymbol{r}(\boldsymbol{x},\boldsymbol{y}) \in \mathcal{C}\,.$$

The necessity of this condition is clear: if there exists $\boldsymbol{y}_0 \in \Delta(\mathcal{B})$ such that

$$R(\boldsymbol{y}_0) \cap \mathcal{C} = \emptyset\,, \qquad \text{where} \qquad R(\boldsymbol{y}_0) = \big\{ \boldsymbol{r}(\boldsymbol{x},\boldsymbol{y}_0),\ \boldsymbol{x} \in \Delta(\mathcal{A}) \big\}$$

then the distance of the compact set $R(\boldsymbol{y}_0)$ to the closed set $\mathcal{C}$ is larger than some $\delta > 0$, while

$$\overline{\boldsymbol{r}}_T = \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{r}(\boldsymbol{x}_t,\boldsymbol{y}_0) = \boldsymbol{r}\big(\overline{\boldsymbol{x}}_T,\, \boldsymbol{y}_0\big) \in R(\boldsymbol{y}_0)$$

for all T, where $\overline{\boldsymbol{x}}_T = \dfrac{1}{T}\displaystyle\sum_{t=1}^{T} \boldsymbol{x}_t$; hence, $d\big(\overline{\boldsymbol{r}}_T,\, \mathcal{C}\big) \geqslant \delta > 0$.

Theorem (Blackwell '56)

A closed convex set $\mathcal{C}$ is approachable if and only

$$\forall\, \boldsymbol{y} \in \Delta(\mathcal{B}), \quad \exists\, \boldsymbol{x} \in \Delta(\mathcal{A}), \qquad \boldsymbol{r}(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{C}\,.$$

The sufficiency of the condition was proved in a constructive way.

Denoting by M a bound in ℓ^2–norm over $\boldsymbol{r}$, i.e.,

$$\max_{(a,b)\in\mathcal{A}\times\mathcal{B}} \big\| \boldsymbol{r}(a,b) \big\|_2 \leqslant M\,,$$

Blackwell's approachability strategy (described on the next slide) ensures that for all strategies of Nature,

$$d\big(\overline{\boldsymbol{r}}_T, \mathcal{C}\big) = \inf_{c\in\mathcal{C}} \left\| c - \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{r}\big(\boldsymbol{x}_t, \boldsymbol{y}_t\big) \right\|_2 \leqslant \frac{2M}{\sqrt{T}}\,.$$

Blackwell's approachability strategy for a closed convex set

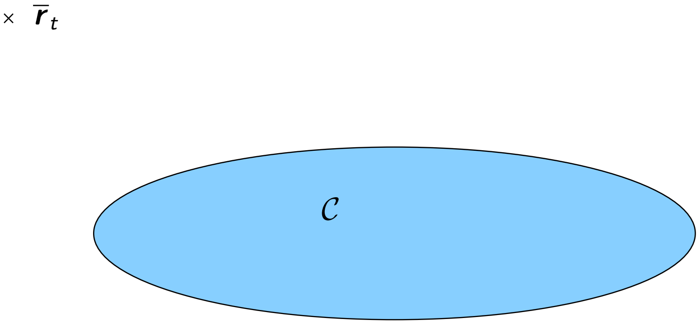

At round t the average payoff is $\overline{\boldsymbol{r}}_t = \frac{1}{t}\sum_{s=1}^{t} \boldsymbol{r}(\boldsymbol{x}_s, \boldsymbol{y}_s)$.

Blackwell's approachability strategy for a closed convex set

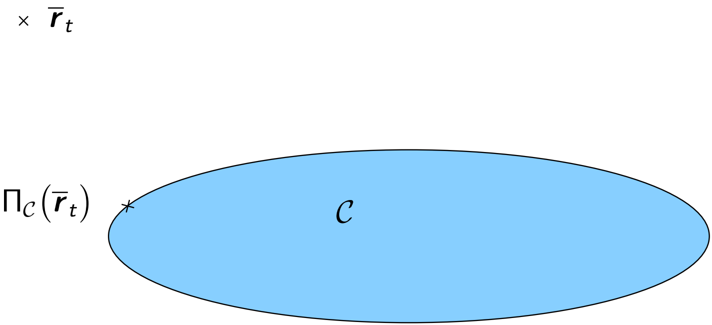

Aim: That $\overline{\boldsymbol{r}}_{t+1}$ gets closer to $\Pi_{\mathcal{C}}(\overline{\boldsymbol{r}}_t)$, the projection of $\overline{\boldsymbol{r}}_t$ onto $\mathcal{C}$.

Blackwell's approachability strategy for a closed convex set

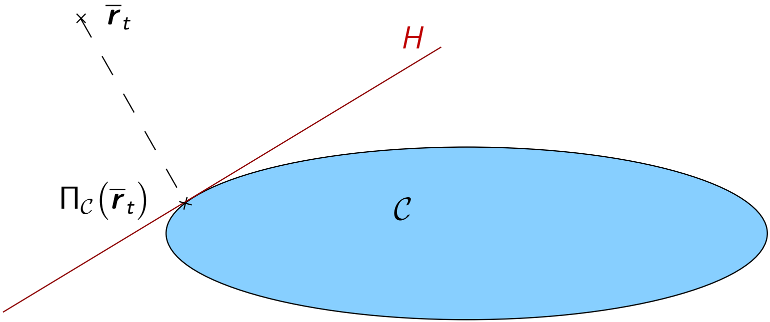

This is true as soon as $r(\boldsymbol{x}_{t+1}, \boldsymbol{y}_{t+1})$ is on the other side of H.

Blackwell's approachability strategy for a closed convex set

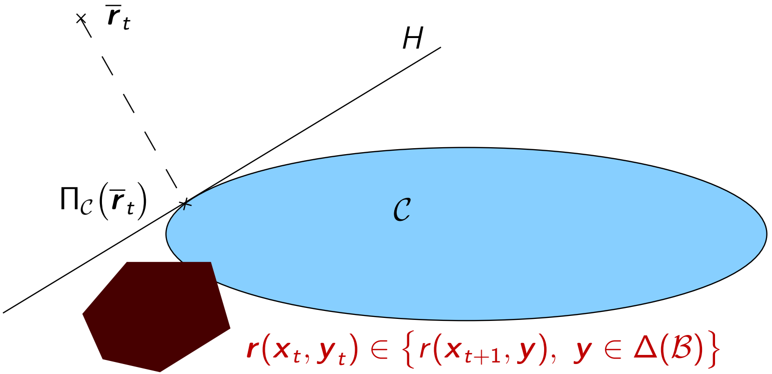

Given H, the existence of a $\boldsymbol{x}_{t+1} \in \Delta(\mathcal{A})$ such that the property illustrated above takes place is guaranteed by the approachability condition (and the minmax theorem); it can be found by solving a minmax program.

An example is formed by the minimization of regret.

A (scalar) payoff function $r : \mathcal{A} \times \mathcal{B} \to \mathbb{R}$ is given and the goal of the decision maker is to ensure that

$$\limsup_{T\to\infty} \frac{1}{T}\left(\sum_{t=1}^{T} r(\boldsymbol{x}_t, \boldsymbol{y}_t) - \max_{a\in\mathcal{A}} \sum_{t=1}^{T} r(a, \boldsymbol{y}_t)\right) \geqslant 0$$

That is, his payoff should on average be almost as large as what he would have obtained by playing a constant action $a \in \mathcal{A}$, all things being equal.

To do so, it suffices to $\boldsymbol{r}$–approach $\mathcal{C}$, where $\boldsymbol{r}(\boldsymbol{x}, \boldsymbol{y}) = \begin{bmatrix} r(\boldsymbol{x}, \boldsymbol{y}) \\ \boldsymbol{y} \end{bmatrix}$

and $\mathcal{C} = \left\{(z, \boldsymbol{y}) \in \mathbb{R} \times \Delta(\mathcal{B}) : \quad z \geqslant \max_{a\in\mathcal{A}} r(a, \boldsymbol{y})\right\}$

Approachability in games with partial monitoring

Our initial motivation and starting point

In games with partial monitoring, the decision maker is not informed of Nature's mixed action $\boldsymbol{y}_t$ at the end of the round t but gets less information.

Description

A finite set of signals $\mathcal{S}$ is available and a matrix $H : \mathcal{A} \times \mathcal{B} \to \Delta(\mathcal{S})$ indicates the feedback received by the decision maker:

He only gets to see $H(\boldsymbol{x}_t, \boldsymbol{y}_t)$ instead of $\boldsymbol{y}_t$ or of $\boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t)$, observation of which would be enough to use Blackwell's strategy.

Yet, he still aims at controlling the average payoff

$$\overline{\boldsymbol{r}}_T = \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{r}(\boldsymbol{x}_t, \boldsymbol{y}_t)$$

Approachability is defined in the same manner as before.

Two mixed actions $\boldsymbol{y}', \boldsymbol{y} \in \Delta(\mathcal{B})$ of Nature are indistinguishable to the decision maker if

$$H(\,\cdot\,,\boldsymbol{y}) = \bigl[H(a,\boldsymbol{y})\bigr]_{a\in\mathcal{A}} = \bigl[H(a,\boldsymbol{y}')\bigr]_{a\in\mathcal{A}} = H(\,\cdot\,,\boldsymbol{y}')$$

This is why we introduced the set-valued mapping

$$\boldsymbol{m} :\ (\boldsymbol{x},\boldsymbol{y}) \longmapsto \Bigl\{\boldsymbol{r}(\boldsymbol{x},\boldsymbol{y}'), \quad \boldsymbol{y}' \ \text{ s.t. } \ H(\,\cdot\,,\boldsymbol{y}') = H(\,\cdot\,,\boldsymbol{y})\Bigr\}$$

There are uncertainties in the obtained payoffs: at best, the decision maker knows that

$$\boldsymbol{r}(\boldsymbol{x}_t,\boldsymbol{y}_t) \ \in \ \boldsymbol{m}(\boldsymbol{x}_t,\boldsymbol{y}_t)$$

[Note: "at best" as he only sees $H(\boldsymbol{x}_t,\boldsymbol{y}_t)$ and not $H(\,\cdot\,,\boldsymbol{y}_t)$, but this is a detail...]

Perchet provided a constructive proof of sufficiency, but for a strategy with an exponentially increasing computational cost (and relying on several layers of notions –internal regret, calibration– all known to be directly related to Blackwell's approachability).

Theorem (Perchet '11)

A closed convex set $\mathcal{C}$ is $\boldsymbol{r}$–approachable with the feedback matrix H if and only

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

Our motivation was mostly to purify the argument but also to exhibit a strategy with constant per-round complexity.

Our hope was that this had already been possible for external regret in this setting (compare Rustichini '99 to Lugosi, Mannor and Stoltz '08)...

The proposed strategy is such that

$$\frac{1}{T}\sum_{t=1}^{T} \boldsymbol{m}(\boldsymbol{x}_t, \boldsymbol{y}_t)$$

converges to $\mathcal{C}$, in the sense that it is eventually included in any ε–neighborhood of $\mathcal{C}$, for $\varepsilon > 0$.

This is of course enough to guarantee that the true average of payoffs $\overline{\boldsymbol{r}}_T$ converges to $\mathcal{C}$.

We therefore introduced the notion of robust approachability, i.e., approachability for set-valued mappings, which, for simplicity, we will assume to be linear. (It is almost the case for the $\boldsymbol{m}$ considered here.)

Robust approachability

The key concept to approachability in games with partial monitoring

A payoff function $\boldsymbol{m}$ associates with each $(a, b) \in \mathcal{A} \times \mathcal{B}$ a subset $\boldsymbol{m}(a, b) \subset \mathbb{R}^d$.

It is linearly extended into a mapping $\boldsymbol{m}$ defined on $\Delta(\mathcal{A} \times \mathcal{B})$.

Definition

A set $\mathcal{C} \subset \mathbb{R}^d$ is $\boldsymbol{m}$–approachable if there exists a strategy for the decision maker such that for all strategies of Nature,

$$\overline{\boldsymbol{m}}_T = \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{m}(\boldsymbol{x}_t, \boldsymbol{y}_t) \subseteq \mathcal{C}_{\varepsilon_T}$$

where $\mathcal{C}_{\varepsilon_T}$ is the ε_T–neighborhood of $\mathcal{C}$, with $\varepsilon_T \to 0$.

Blackwell's strategy to the farthest point does not work

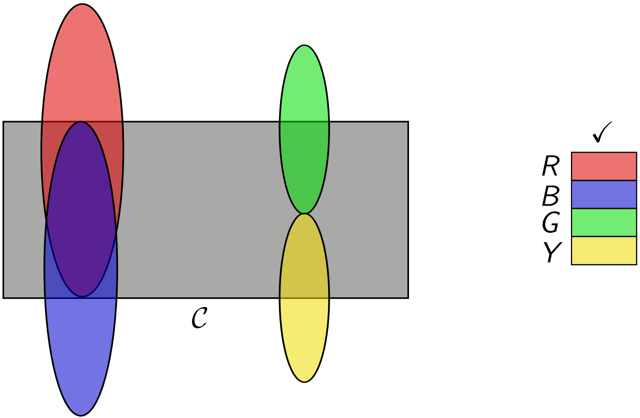

Actions sets are $\mathcal{A} = \{R, B, G, Y\}$ and $\mathcal{B} = \{\checkmark\}$

$\boldsymbol{m}(R, \checkmark)$ is the red set, $\boldsymbol{m}(B, \checkmark)$ is the blue set, and so on

Blackwell's strategy to the farthest point does not work

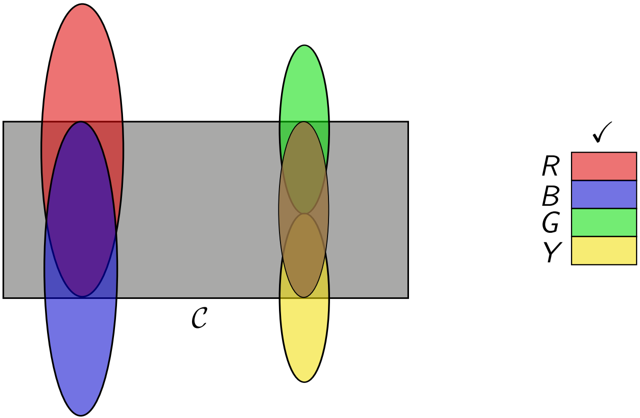

In this case, the stated condition is satisfied, with $x = \frac{1}{2}\delta_G + \frac{1}{2}\delta_Y$

Blackwell's strategy to the farthest point does not work

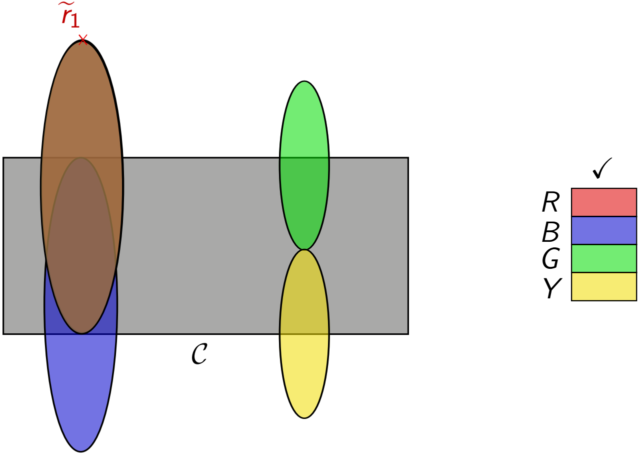

Assume that R was played; the set $\overline{\boldsymbol{m}}_1$ is in brown

The farthest point in $\overline{\boldsymbol{m}}_1$ to $\mathcal{C}$ is denoted by $\widetilde{r}_1$

Blackwell's strategy to the farthest point does not work

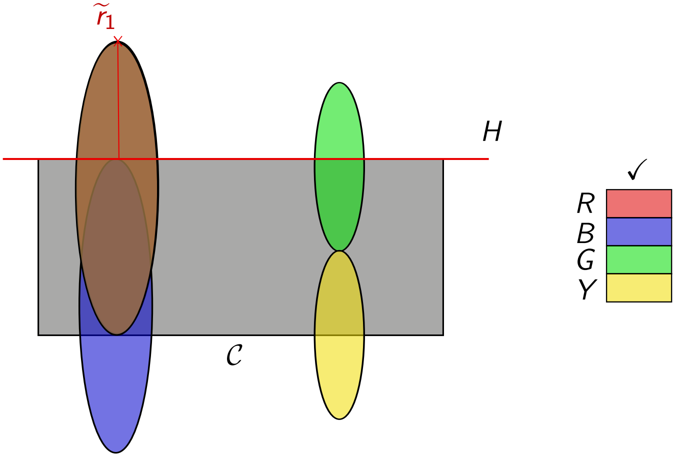

The blue set is on the other side of the hyperplane

Hence, playing B at stage 2 is a choice compatible with Blackwell's strategy

Blackwell's strategy to the farthest point does not work

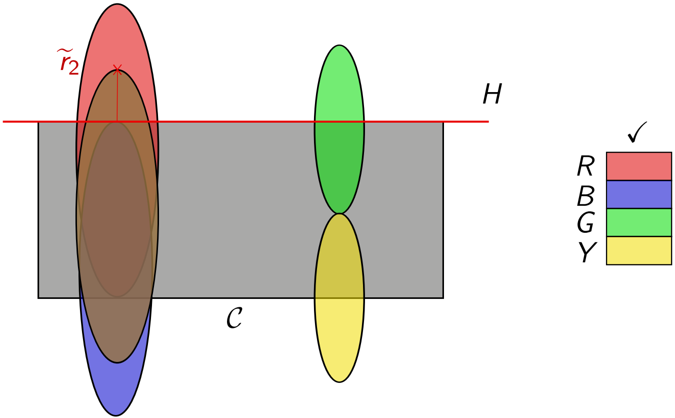

The blue set is still on the other side of the hyperplane

Hence, playing B at stage 3 is still a choice compatible with Blackwell's strategy

Blackwell's strategy to the farthest point does not work

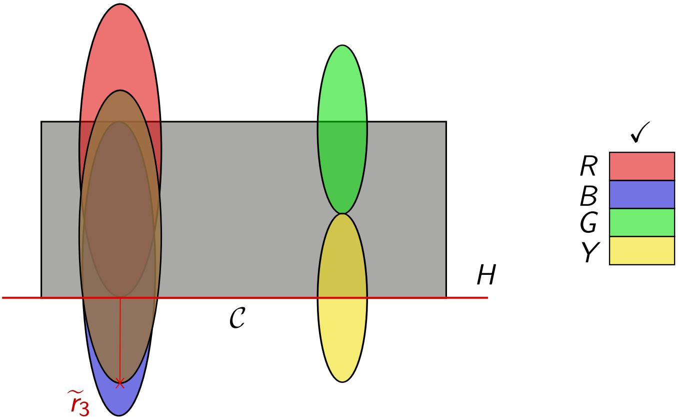

Now, the red set is on the other side of the hyperplane

Hence, playing R at stage 3 is a choice compatible with Blackwell's strategy

Blackwell's strategy to the farthest point does not work

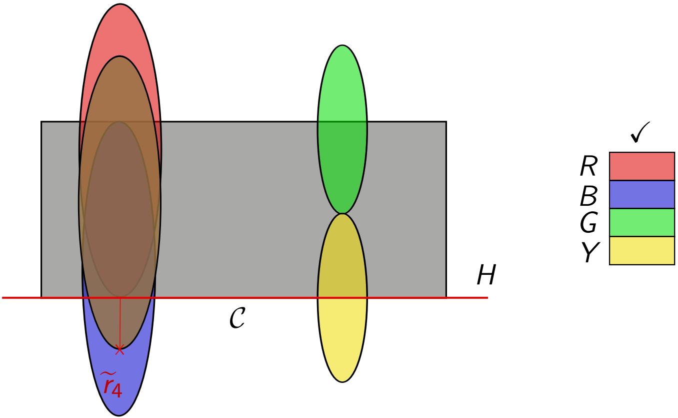

The algorithm can oscillate indefinitely between R and B without $\overline{\boldsymbol{m}}_T$ converging to $\mathcal{C}$

The payoff sets must be considered in their entirety!

An example of a good path of actions

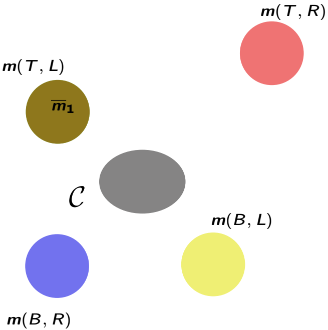

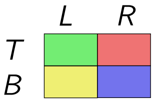

Sequence played:

(T, L)

Action sets $\mathcal{A} = \{T, B\}$ and $\mathcal{B} = \{L, R\}$; the set $\mathcal{C}$ is $\boldsymbol{m}$–approachable

If T and L are chosen at the first round: $\overline{\boldsymbol{m}}_1 = \boldsymbol{m}(T, L)$ is in brown

An example of a good path of actions

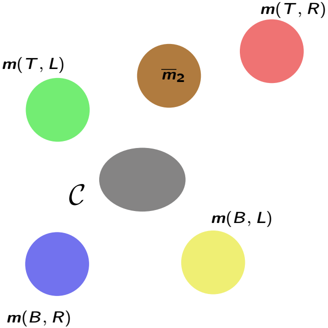

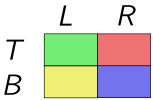

Sequence played:

$(T, L); (T, R)$

Action sets $\mathcal{A} = \{T, B\}$ and $\mathcal{B} = \{L, R\}$; the set $\mathcal{C}$ is $\boldsymbol{m}$–approachable

Second round: (T, R) is played, $\overline{\boldsymbol{m}}_2 = \frac{1}{2}\boldsymbol{m}(T, L) + \frac{1}{2}\boldsymbol{m}(T, R)$

An example of a good path of actions

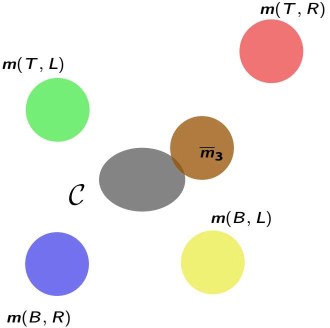

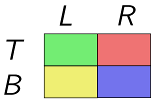

Sequence played:

$(T, L); (T, R); (B, L)$

Action sets $\mathcal{A} = \{T, B\}$ and $\mathcal{B} = \{L, R\}$; the set $\mathcal{C}$ is $\boldsymbol{m}$–approachable

Third round: $\overline{\boldsymbol{m}}_3 = \frac{1}{3}\boldsymbol{m}(T, L) + \frac{1}{3}\boldsymbol{m}(T, R) + \frac{1}{3}\boldsymbol{m}(B, L)$

An example of a good path of actions

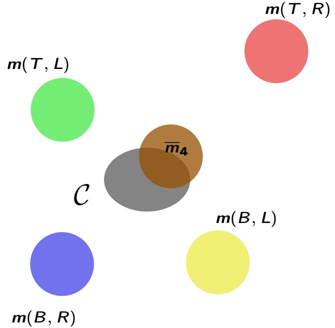

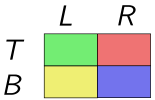

Sequence played:

$(T, L); (T, R); (B, L); (B, R)$

Action sets $\mathcal{A} = \{T, B\}$ and $\mathcal{B} = \{L, R\}$; the set $\mathcal{C}$ is $\boldsymbol{m}$–approachable

$$\overline{\boldsymbol{m}}_4 = \frac{1}{4}\boldsymbol{m}(T, L) + \frac{1}{4}\boldsymbol{m}(T, R) + \frac{1}{4}\boldsymbol{m}(B, L) + \frac{1}{4}\boldsymbol{m}(B, R)$$

An example of a good path of actions

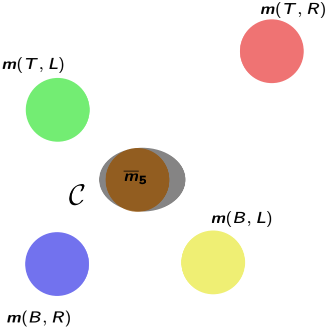

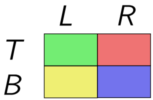

Sequence played:

$(T, L); (T, R); (B, L); (B, R); (B, R)$

Action sets $\mathcal{A} = \{T, B\}$ and $\mathcal{B} = \{L, R\}$; the set $\mathcal{C}$ is $\boldsymbol{m}$–approachable

$$\overline{\boldsymbol{m}}_5 = \frac{1}{5}\boldsymbol{m}(T, L) + \frac{1}{5}\boldsymbol{m}(T, R) + \frac{1}{5}\boldsymbol{m}(B, L) + \frac{2}{5}\boldsymbol{m}(B, R)$$

An example of a good path of actions

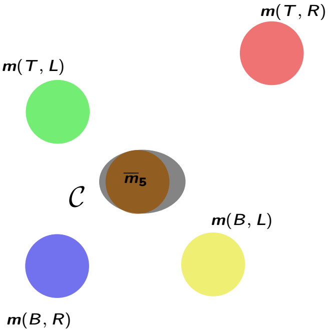

	L	R
T		
B		

Sequence played:

(T, L); (T, R); (B, L); (B, R); (B, R)

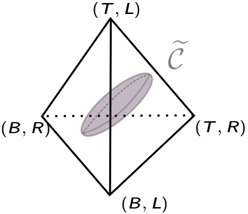

$$\frac{1}{5}\boldsymbol{m}(T,L) + \frac{1}{5}\boldsymbol{m}(T,R) + \frac{1}{5}\boldsymbol{m}(B,L) + \frac{2}{5}\boldsymbol{m}(B,R) \subseteq \mathcal{C}, \text{ that is,}$$

$$\left(\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{2}{5}\right) \in \widetilde{\mathcal{C}} = \left\{\mu \in \Delta(\mathcal{A} \times \mathcal{B}) : \quad \mathbb{E}_\mu\left[\boldsymbol{m}(A,B)\right] \subseteq \mathcal{C}\right\}$$

An example of a good path of actions

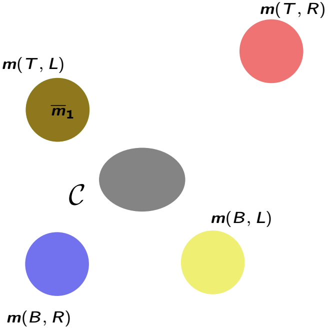

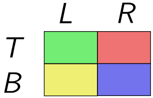

Sequence played:

(T, L)

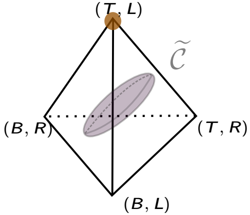

Approaching $\mathcal{C} \subset \mathbb{R}^d$ is equivalent to approaching $\widetilde{\mathcal{C}} \subset \Delta(\mathcal{A} \times \mathcal{B})$

Abstract payoff of first round: $\boldsymbol{a}_1 = \delta_{(T,L)}$

An example of a good path of actions

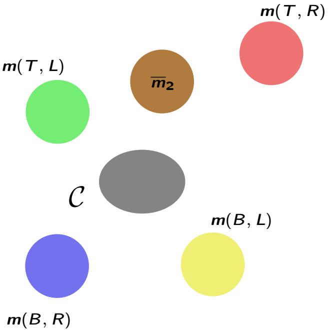

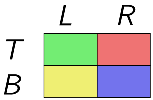

Sequence played:

$(T, L); (T, R)$

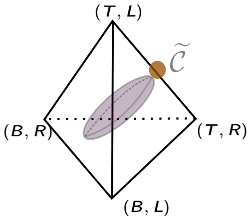

Approaching $\mathcal{C} \subset \mathbb{R}^d$ is equivalent to approaching $\widetilde{\mathcal{C}} \subset \Delta(\mathcal{A} \times \mathcal{B})$

Average payoff after second round: $\overline{a}_2 = \frac{1}{2}\delta_{(T,L)} + \frac{1}{2}\delta_{(T,R)}$

An example of a good path of actions

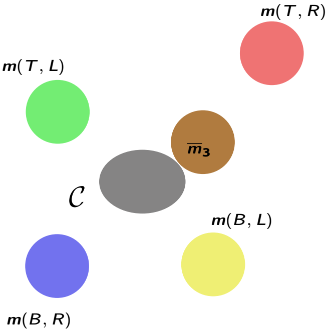

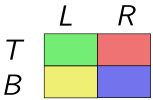

Sequence played:

$(T, L); (T, R); (B, L)$

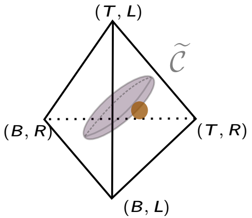

Approaching $\mathcal{C} \subset \mathbb{R}^d$ is equivalent to approaching $\widetilde{\mathcal{C}} \subset \Delta(\mathcal{A} \times \mathcal{B})$

After third round: $\overline{\boldsymbol{a}}_3 = \frac{1}{3}\delta_{(T,L)} + \frac{1}{3}\delta_{(T,R)} + \frac{1}{3}\delta_{(B,L)}$

An example of a good path of actions

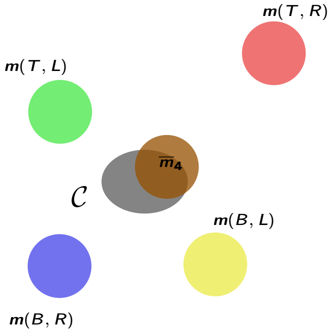

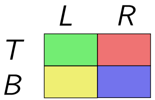

Sequence played:

$(T, L); (T, R); (B, L); (B, R)$

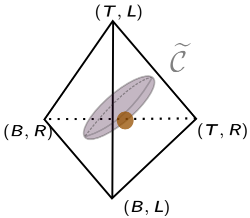

Approaching $\mathcal{C} \subset \mathbb{R}^d$ is equivalent to approaching $\widetilde{\mathcal{C}} \subset \Delta(\mathcal{A} \times \mathcal{B})$

$$\overline{\boldsymbol{a}}_4 = \frac{1}{4}\delta_{(T,L)} + \frac{1}{4}\delta_{(T,R)} + \frac{1}{4}\delta_{(B,L)} + \frac{1}{4}\delta_{(B,R)}$$

An example of a good path of actions

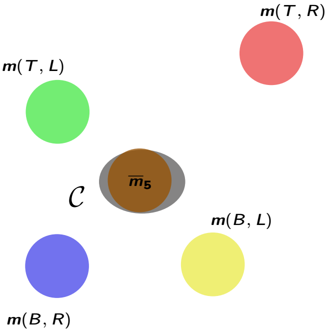

	L	R
T		
B		

Sequence played:

$(T, L); (T, R); (B, L); (B, R); (B, R)$

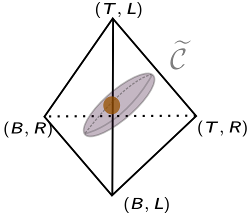

Approaching $\mathcal{C} \subset \mathbb{R}^d$ is equivalent to approaching $\widetilde{\mathcal{C}} \subset \Delta(\mathcal{A} \times \mathcal{B})$

$$\overline{\boldsymbol{a}}_5 = \frac{1}{5}\delta_{(T,L)} + \frac{1}{5}\delta_{(T,R)} + \frac{1}{5}\delta_{(B,L)} + \frac{2}{5}\delta_{(B,R)}$$

There is an equivalence between the following two settings.

Payoffs with uncertainties

- Actions taken in $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$
- Payoff given by the subset $\boldsymbol{m}(\boldsymbol{x},\boldsymbol{y}) \subset \mathbb{R}^d$
- Target closed convex set $\mathcal{C}$

Payoffs without uncertainties (classical setting)

- Actions taken in $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$
- Payoff given by a single point: the product-distribution $\boldsymbol{x} \otimes \boldsymbol{y} \in \mathbb{R}^{\mathcal{A}\times\mathcal{B}}$
- Target closed convex set $\widetilde{\mathcal{C}}$, where

$$\widetilde{\mathcal{C}} = \Big\{\mu \in \Delta(\mathcal{A}\times\mathcal{B}) : \quad \mathbb{E}_\mu\big[\boldsymbol{m}(A,B)\big] \subset \mathcal{C}\Big\}$$

In the sense that $\mathcal{C}$ is $\boldsymbol{m}$–approachable
if and only if $\widetilde{\mathcal{C}}$ is $\otimes$–approachable

This equivalence gives rise to an auxiliary strategy based on which we construct a strategy to perform approachability in the context of games with partial monitoring.

Instead of digging into the details, I just mention some of the tools needed to complete the proof:

- $\boldsymbol{m}$ is only piecewise linear in partial monitoring (under some mild conditions on r or on $\mathcal{C}$) but is induced by a linear function $\widetilde{\boldsymbol{m}}$ defined on a lifted space;
- we need to play in blocks as we do not observe the $H(\,\cdot\,, \boldsymbol{y}_t)$ but only the $H(\boldsymbol{x}_t, \boldsymbol{y}_t)$;
- some exploration–exploitation tradeoff is useful as well...

Note: Convergence rates (of the suboptimal order of $T^{-1/5}$) can be obtained.

This equivalence gives rise to an auxiliary strategy based on which we construct a strategy to perform approachability in the context of games with partial monitoring.

But that is enough for today...

Thanks for your attention...

And thanks for staying till the end of this workshop!