

Alternation makes the adversary weaker

Stratis Skoulakis

EPFL



joint work with



V. Cevher, **EPFL**



A. Cutkosky, **BU**



G. Piliouras,
DeepMind




A. Kavis, **UT Austin**



L. Viano, **EPFL**

Normal Form Games

Rock-Paper-Scissors



Payoff matrix of Alice


Alice/Bob	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

Payoff matrix of Bob

Alice/Bob	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

Normal Form Games

Rock-Paper-Scissors



Alice plays Rock

Alice/Bob	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

Alice gets -1

Bob plays Scissors

Alice/Bob	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

Bob gets 1

Normal Form Games

Normal-Form Games

Matrices A and B

Payoff matrix A of Alice

A/B	1	...	j	...	m
1	A_{11}	...	A_{1j}	...	A_{1m}
\vdots	\vdots	\vdots	...	\vdots	\vdots
i	A_{i1}	...	A_{ij}	...	A_{im}
\vdots	\vdots	\vdots	...	\vdots	\vdots
n	A_{n1}	...	A_{nj}	...	A_{nm}

Payoff matrix B of Bob

A/B	1	...	j	...	m
1	B_{11}	...	B_{1j}	...	B_{1m}
\vdots	\vdots	\vdots	...	\vdots	\vdots
i	B_{i1}	...	B_{ij}	...	B_{im}
\vdots	\vdots	\vdots	...	\vdots	\vdots
n	B_{n1}	...	B_{nj}	...	B_{nm}

Normal Form Games

Normal-Form Games

Matrices A and B

Alice plays action i

A/B	1	...	j	...	m
1	A_{11}	...	A_{1j}	...	A_{1m}
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
i	A_{i1}	...	A_{ij}	...	A_{im}
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	A_{n1}	...	A_{nj}	...	A_{nm}

Alice gets A_{ij}

Bob plays action j

A/B	1	...	j	...	m
1	B_{11}	...	B_{1j}	...	B_{1m}
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
i	B_{i1}	...	B_{ij}	...	B_{im}
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	B_{n1}	...	B_{nj}	...	B_{nm}

Bob gets B_{ij}

Normal Form Games

Normal-Form Games - Mixed Strategies

Matrices A and B

Alice plays a prob. distr.

$$x = (x_1, \dots, x_n) \in \Delta_n$$

A/B	1	...	j	...	m
1	A_{11}	...	A_{1j}	...	A_{1m}
⋮	⋮	⋮	...	⋮	⋮
i	A_{i1}	...	A_{ij}	...	A_{im}
⋮	⋮	⋮	...	⋮	⋮
n	A_{n1}	...	A_{nj}	...	A_{nm}

Bob plays a prob. distr.

$$y = (y_1, \dots, y_m) \in \Delta_m$$

A/B	1	...	j	...	m
1	B_{11}	...	B_{1j}	...	B_{1m}
⋮	⋮	⋮	...	⋮	⋮
i	B_{i1}	...	B_{ij}	...	B_{im}
⋮	⋮	⋮	...	⋮	⋮
n	B_{n1}	...	B_{nj}	...	B_{nm}

Normal Form Games

Normal-Form Games - Mixed Strategies

Matrices A and B

Alice plays a prob. distr.

$$x = (x_1, \dots, x_n) \in \Delta_n$$

A/B	1	...	j	...	m
1	A_{11}	...	A_{1j}	...	A_{1m}
\vdots	\vdots	\vdots	...	\vdots	\vdots
i	A_{i1}	...	A_{ij}	...	A_{im}
\vdots	\vdots	\vdots	...	\vdots	\vdots
n	A_{n1}	...	A_{nj}	...	A_{nm}

Alice's expected cost

$$x^\top Ay$$

Bob plays a prob. distr.

$$y = (y_1, \dots, y_m) \in \Delta_m$$

A/B	1	...	j	...	m
1	B_{11}	...	B_{1j}	...	B_{1m}
\vdots	\vdots	\vdots	...	\vdots	\vdots
i	B_{i1}	...	B_{ij}	...	B_{im}
\vdots	\vdots	\vdots	...	\vdots	\vdots
n	B_{n1}	...	B_{nj}	...	B_{nm}

Bob's expected cost

$$x^\top By$$

Normal Form Games

Normal-Form Games - Mixed Strategies

Matrices A and B

Alice plays a prob. distr.

$$x = (x_1, \dots, x_n) \in \Delta_n$$

A/B	1	...	j	...	m
1	A_{11}	...	A_{1j}	...	A_{1m}
\vdots	\vdots	\vdots	...	\vdots	\vdots
i	A_{i1}	...	A_{ij}	...	A_{im}
\vdots	\vdots	\vdots	...	\vdots	\vdots
n	A_{n1}	...	A_{nj}	...	A_{nm}

$$x^\top A y$$

Bob plays a prob. distr.

$$y = (y_1, \dots, y_m) \in \Delta_m$$

A/B	1	...	j	...	m
1	B_{11}	...	B_{1j}	...	B_{1m}
\vdots	\vdots	\vdots	...	\vdots	\vdots
i	B_{i1}	...	B_{ij}	...	B_{im}
\vdots	\vdots	\vdots	...	\vdots	\vdots
n	B_{n1}	...	B_{nj}	...	B_{nm}

$$x^\top B y$$

Normal-Form Games over Time

- Alice and Bob play the normal-form game (A, B) over T rounds.
 - ▶ *Simultaneous play*: Agents simultaneously select their strategies at each round (*very well-studied*).
 - ▶ *Alternating play*: Agents alternatingly update their strategies (*this work*).

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects y_1

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects y_2

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \longrightarrow$ Alice gets $x_2^\top A y_2$

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \longrightarrow$ Alice gets $x_2^\top A y_2$
- Alice selects x_3 / Bob selects y_3

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \longrightarrow$ Alice gets $x_2^\top A y_2$
- Alice selects x_3 / Bob selects $y_3 \longrightarrow$ Alice gets $x_3^\top A y_3$
- ...

Simultaneous Play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \longrightarrow$ Alice gets $x_2^\top A y_2$
- Alice selects x_3 / Bob selects $y_3 \longrightarrow$ Alice gets $x_3^\top A y_3$
- ...

How should Alice select her actions over time?

Simultaneous Play

Simultaneous Play

Bob challenges **Alice** to play (A, B) for T rounds.

- **Alice** selects x_1 / **Bob** selects $y_1 \rightarrow$ **Alice** gets $x_1^\top A y_1$
- **Alice** selects x_2 / **Bob** selects $y_2 \rightarrow$ **Alice** gets $x_2^\top A y_2$
- **Alice** selects x_3 / **Bob** selects $y_3 \rightarrow$ **Alice** gets $x_3^\top A y_3$
- ...

How should **Alice** select her actions over time? *No Regret algorithms*

- Regret Matching [Blackwell '65]
- Hedge [Freund et al. '97]
- Online Gradient Descent [Zinkevich '03]
- Follow the Regularized Leader [Abernethy et al. '10]
-

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \longrightarrow$ Alice gets $x_2^\top A y_2$
- ...

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \rightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \rightarrow$ Alice gets $x_2^\top A y_2$
- ...

Hedge Algorithm [Freund and Schapire '97] Godel Prize '03

$$x_{t+1}(i) = \frac{x_t(i) \cdot e^{-\gamma[Ay_t]_i}}{\sum_{j=1}^n x_t(j) \cdot e^{-\gamma[Ay_t]_j}}$$

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \rightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \rightarrow$ Alice gets $x_2^\top A y_2$
- ...

Hedge Algorithm [Freund and Schapire '97] Godel Prize '03

$$x_{t+1}(i) = \frac{x_t(i) \cdot e^{-\gamma[Ay_t]_i}}{\sum_{j=1}^n x_t(j) \cdot e^{-\gamma[Ay_t]_j}}$$

Theorem (Freund and Schapire JCSS '97)

No matter Bob's strategies y_1, \dots, y_T , the regret of Alice $\mathcal{R}(T)$

$$\mathcal{R}(T) := \underbrace{\sum_{t=1}^T x_t^\top A y_t}_{\text{cost of Alice}} - \underbrace{\min_{i \in [n]} \sum_{t=1}^T [Ay_t]_i}_{\text{cost of best action}} \leq \tilde{O}(\sqrt{T})$$

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \longrightarrow$ Alice gets $x_2^\top A y_2$
- Alice selects x_3 / Bob selects $y_3 \longrightarrow$ Alice gets $x_3^\top A y_3$
- ...

Theorem (Freund and Schapire JCSS '97)

No matter Bob's strategies,

$$\underbrace{\frac{1}{T} \sum_{t=1}^T x_t^\top A y_t}_{\text{time-average cost}} \leq \underbrace{\frac{1}{T} \min_{x \in \Delta_n} \sum_{t=1}^T x^\top A y_t}_{\text{best fixed action}} + \tilde{O}\left(\frac{\sqrt{T}}{T}\right) \rightarrow 0$$

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \rightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \rightarrow$ Alice gets $x_2^\top A y_2$
- ...

Theorem (Freund and Schapire JCSS '97)

No matter Bob's strategies y_1, \dots, y_T , the regret of Alice $\mathcal{R}(T)$

$$\mathcal{R}(T) := \underbrace{\sum_{t=1}^T x_t^\top A y_t}_{\text{cost of Alice}} - \underbrace{\min_{i \in [n]} \sum_{t=1}^T [A y_t]_i}_{\text{cost of best action}} \leq \tilde{O}(\sqrt{T})$$

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \rightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \rightarrow$ Alice gets $x_2^\top A y_2$
- ...

Theorem (Freund and Schapire JCSS '97)

No matter Bob's strategies y_1, \dots, y_T , the regret of Alice $\mathcal{R}(T)$

$$\mathcal{R}(T) := \underbrace{\sum_{t=1}^T x_t^\top A y_t}_{\text{cost of Alice}} - \underbrace{\min_{i \in [n]} \sum_{t=1}^T [A y_t]_i}_{\text{cost of best action}} \leq \tilde{O}(\sqrt{T})$$

Can Alice do better?

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \rightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \rightarrow$ Alice gets $x_2^\top A y_2$
- ...

Theorem (Freund and Schapire JCSS '97)

No matter Bob's strategies y_1, \dots, y_T , the regret of Alice $\mathcal{R}(T)$

$$\mathcal{R}(T) := \underbrace{\sum_{t=1}^T x_t^\top A y_t}_{\text{cost of Alice}} - \underbrace{\min_{i \in [n]} \sum_{t=1}^T [A y_t]_i}_{\text{cost of best action}} \leq \tilde{O}(\sqrt{T})$$

Can Alice do better? **No!**

Simultaneous play

Simultaneous Play

Bob challenges **Alice** to play (A, B) for T rounds.

- **Alice** selects x_1 / **Bob** selects $y_1 \longrightarrow$ **Alice** gets $x_1^\top A y_1$
- **Alice** selects x_2 / **Bob** selects $y_2 \longrightarrow$ **Alice** gets $x_2^\top A y_2$
- ...

Theorem (folklore)

Bob can always select y_1, \dots, y_T , **the regret of Alice**

$$\mathcal{R}(T) := \underbrace{\sum_{t=1}^T x_t^\top A y_t}_{\text{cost of Alice}} - \underbrace{\min_{i \in [n]} \sum_{t=1}^T [A y_t]_i}_{\text{cost of best action}} \geq \Omega(\sqrt{T})$$

$\mathcal{O}(\sqrt{T})$ regret is the best **Alice** can get in **simultaneous play!**

Simultaneous play

Simultaneous Play

Bob challenges **Alice** to play (A, B) for T rounds.

- **Alice** selects x_1 / **Bob** selects $y_1 \longrightarrow$ **Alice** gets $x_1^\top A y_1$
- **Alice** selects x_2 / **Bob** selects $y_2 \longrightarrow$ **Alice** gets $x_2^\top A y_2$
- ...

Theorem (folklore)

Bob can always select y_1, \dots, y_T , *the regret of Alice*

$$\mathcal{R}(T) := \underbrace{\sum_{t=1}^T x_t^\top A y_t}_{\text{cost of Alice}} - \underbrace{\min_{i \in [n]} \sum_{t=1}^T [A y_t]_i}_{\text{cost of best action}} \geq \Omega(\sqrt{T})$$

Simultaneous play

Simultaneous Play

Bob challenges Alice to play (A, B) for T rounds.

- Alice selects x_1 / Bob selects $y_1 \longrightarrow$ Alice gets $x_1^\top A y_1$
- Alice selects x_2 / Bob selects $y_2 \longrightarrow$ Alice gets $x_2^\top A y_2$
- ...

Theorem (folklore)

Bob can always select y_1, \dots, y_T , the regret of Alice

$$\mathcal{R}(T) := \underbrace{\sum_{t=1}^T x_t^\top A y_t}_{\text{cost of Alice}} - \underbrace{\min_{i \in [n]} \sum_{t=1}^T [A y_t]_i}_{\text{cost of best action}} \geq \Omega(\sqrt{T})$$

What if Alice and Bob play in alternating turns?

Solving Heads'up Poker



Online Learning + Simultaneous Play

Polaris [Bowling et al. AAMAS '09] → Decent Performance

Solving Heads'up Poker



Online Learning + Simultaneous Play

Polaris [Bowling et al. AAMAS '09] → Decent Performance

Online Learning + Alternating Play

Cepheus [Oskari et al. IJCAI '15], Libratus [Brown et al. IJCAI '17] → **Beat Human Experts!**

Solving Heads'up Poker



Online Learning + Simultaneous Play

Polaris [Bowling et al. AAMAS '09] → Decent Performance

Online Learning + Alternating Play

Cepheus [Oskari et al. IJCAI '15], Libratus [Brown et al. IJCAI '17] → **Beat Human Experts!**

Regret Guarantees → Faster Training!

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects x_1

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \longrightarrow$ Alice gets $x_1^\top A y_0$

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \longrightarrow$ Alice gets $x_1^\top A y_0$
- Bob selects y_2

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \rightarrow$ Alice gets $x_1^\top A y_0$
- Bob selects $y_2 \rightarrow$ Alice gets $x_1^\top A y_2$

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \rightarrow$ Alice gets $x_1^\top A y_0$
- Bob selects $y_2 \rightarrow$ Alice gets $x_1^\top A y_2$
- Alice selects x_3

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \longrightarrow$ Alice gets $x_1^\top A y_0$
- Bob selects $y_2 \longrightarrow$ Alice gets $x_1^\top A y_2$
- Alice selects $x_3 \longrightarrow$ Alice gets $x_3^\top A y_2$

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \rightarrow$ Alice gets $x_1^\top A y_0$
- Bob selects $y_2 \rightarrow$ Alice gets $x_1^\top A y_2$
- Alice selects $x_3 \rightarrow$ Alice gets $x_3^\top A y_2$
- ...

Alternating Turns

Alternating Play

Bob challenges **Alice** to play (A, B) for T rounds in **alternating turns**.

- **Bob** selects y_0
- **Alice** selects $x_1 \rightarrow$ **Alice** gets $x_1^\top A y_0$
- **Bob** selects $y_2 \rightarrow$ **Alice** gets $x_1^\top A y_2$
- **Alice** selects $x_3 \rightarrow$ **Alice** gets $x_3^\top A y_2$
- ...

Can **Alice** have regret $\mathcal{R}(T)$ better than $\mathcal{O}(\sqrt{T})$?

$$\mathcal{R}(T) := \underbrace{\sum_{k=0}^{T/2} x_{2k+1}^\top A (y_{2k} + y_{2k+2})}_{\text{Alice's cost}} - \underbrace{\min_{x \in \Delta_n} \sum_{k=0}^{T/2} x^\top A (y_{2k} + y_{2k+2})}_{\text{best fixed action}}$$

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \rightarrow$ Alice gets $x_1^\top A y_0$
- Bob selects $y_2 \rightarrow$ Alice gets $x_1^\top A y_2$
- Alice selects $x_3 \rightarrow$ Alice gets $x_3^\top A y_2$
- ...

Theorem ([Cevher, Cutkovsky*, Piliouras, Skoulakis*, Viano NeurIPS '23 spotlight])

In alternating play, Alice can always guarantee

- $\tilde{\mathcal{O}}(T^{1/3})$ regret for general games.
- $\mathcal{O}(\log T)$ for $n = 2$ actions (different algorithm).

Alternating Turns

Previous Results

- If **both** Alice and Bob use *Gradient Descent* in **unconstrained zero-sum games** $\rightarrow \mathcal{O}(1)$ regret [Bailey et al. COLT 2020]
- If **both** Alice and Bob use *Hedge* in **zero-sum games** $\rightarrow \mathcal{O}(T^{1/3})$ regret [Wibisono et al. NeurIPS 2022]

Alternating Turns

Previous Results

- If **both** Alice and Bob use *Gradient Descent* in **unconstrained zero-sum games** $\rightarrow \mathcal{O}(1)$ regret [Bailey et al. COLT 2020]
- If **both** Alice and Bob use *Hedge* in **zero-sum games** $\rightarrow \mathcal{O}(T^{1/3})$ regret [Wibisono et al. NeurIPS 2022]

Previous results make assumptions on Bob's behavior.

Alternating Turns

Previous Results

- If **both Alice** and **Bob** use *Gradient Descent* in **unconstrained zero-sum games** $\rightarrow \mathcal{O}(1)$ regret [Bailey et al. COLT 2020]
- If **both Alice** and **Bob** use *Hedge* in **zero-sum games** $\rightarrow \mathcal{O}(T^{1/3})$ regret [Wibisono et al. NeurIPS 2022]

Previous results make assumptions on Bob's behavior.

Theorem ([Cevher, Cutkovsky*, Piliouras, Skoulakis*, Viano NeurIPS '23 spotlight])

In alternating play, Alice can always guarantee

- $\tilde{\mathcal{O}}(T^{1/3})$ regret for general games.
- $\mathcal{O}(\log T)$ for $n = 2$ actions (different algorithm).

no matter Bob's behavior!

Alternating Turns

Alternating Play

Bob challenges Alice to play (A, B) for T rounds in **alternating turns**.

- Bob selects y_0
- Alice selects $x_1 \rightarrow$ Alice gets $x_1^\top A y_0$
- Bob selects $y_2 \rightarrow$ Alice gets $x_1^\top A y_2$
- Alice selects $x_3 \rightarrow$ Alice gets $x_3^\top A y_2$
- ...

	Simultaneous Play	Alternating Play
General Games	$\tilde{O}(\sqrt{T})$ tight	$\tilde{O}(T^{1/3})$
2 actions for Alice	$\tilde{O}(\sqrt{T})$ tight	$\tilde{O}(\log T)$

Alternating Turns

Alternating Play

Bob challenges **Alice** to play (A, B) for T rounds in **alternating turns**.

- **Bob** selects y_0
- **Alice** selects $x_1 \rightarrow$ **Alice** gets $x_1^\top A y_0$
- **Bob** selects $y_2 \rightarrow$ **Alice** gets $x_1^\top A y_2$
- **Alice** selects $x_3 \rightarrow$ **Alice** gets $x_3^\top A y_2$
- ...

Our algorithm (general simplex)

$$x_{2k+1} \leftarrow \operatorname{argmin}_{x \in \Delta_n} \left[3 \underbrace{x^\top A y_{2k}}_{\text{exploits } y_{2k}} + 2 \underbrace{\sum_{k'=0}^{k-1} x^\top A y_{2k'}}_{\text{reinforces good past actions}} - \underbrace{\gamma \cdot \sum_{i=1}^n \log x_i}_{\text{prevents overfitting}} \right]$$

Alternating Turns

Alternating Play

Bob challenges **Alice** to play (A, B) for T rounds in **alternating turns**.

- **Bob** selects y_0
- **Alice** selects $x_1 \rightarrow$ **Alice** gets $x_1^\top A y_0$
- **Bob** selects $y_2 \rightarrow$ **Alice** gets $x_1^\top A y_2$
- **Alice** selects $x_3 \rightarrow$ **Alice** gets $x_3^\top A y_2$
- ...

Our algorithm (general simplex)

$$x_{2k+1} \leftarrow \operatorname{argmin}_{x \in \Delta_n} \left[3 \underbrace{x^\top A y_{2k}}_{\text{exploits } y_{2k}} + 2 \underbrace{\sum_{k'=0}^{k-1} x^\top A y_{2k'}}_{\text{reinforces good past actions}} - \underbrace{\gamma \cdot \sum_{i=1}^n \log x_i}_{\text{prevents overfitting}} \right]$$

Setting $\gamma = \mathcal{O}(T^{1/3}) \rightarrow \mathcal{O}(T^{1/3})$ alternating regret

2 action simplex

- If Alice admits 2 actions $\rightarrow \mathcal{O}(\log T)$ in alternating play!

2 action simplex

- If Alice admits 2 actions $\rightarrow \mathcal{O}(\log T)$ in alternating play!
- At round $t = 2k + 1$

2 action simplex

- If Alice admits 2 actions $\rightarrow \mathcal{O}(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - ▶ Greedy Best Response $\rightarrow w_k = \arg \min_{x \in \Delta_n} [x^\top A y_{2k}]$

2 action simplex

- If Alice admits 2 actions $\rightarrow \mathcal{O}(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - ▶ Greedy Best Response $\rightarrow w_k = \arg \min_{x \in \Delta_n} [x^\top A y_{2k}]$
 - ▶ Follow the Regularized Leader \rightarrow

$$z_k = \arg \min_{x \in \Delta_n} \left[\sum_{k=1}^{t/2} x^\top A y_{2k} + \gamma \|x\|^2 \right]$$

2 action simplex

- If Alice admits 2 actions $\rightarrow \mathcal{O}(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - ▶ Greedy Best Response $\rightarrow w_k = \arg \min_{x \in \Delta_n} [x^\top Ay_{2k}]$
 - ▶ Follow the Regularized Leader \rightarrow

$$z_k = \arg \min_{x \in \Delta_n} \left[\sum_{k=1}^{t/2} x^\top Ay_{2k} + \gamma \|x\|^2 \right]$$

- ▶ If FTRL (z_k) admits $\mathcal{O}(\sqrt{T})$ regret \rightarrow Greedy BR (w_k) admits $\mathcal{O}(1)$ regret

2 action simplex

- If Alice admits 2 actions $\rightarrow \mathcal{O}(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - ▶ Greedy Best Response $\rightarrow w_k = \arg \min_{x \in \Delta_n} [x^\top A y_{2k}]$
 - ▶ Follow the Regularized Leader \rightarrow

$$z_k = \arg \min_{x \in \Delta_n} \left[\sum_{k=1}^{t/2} x^\top A y_{2k} + \gamma \|x\|^2 \right]$$

- ▶ If FTRL (z_k) admits $\mathcal{O}(\sqrt{T})$ regret \rightarrow Greedy BR (w_k) admits $\mathcal{O}(1)$ regret
- ▶ If Greedy BR (w_k) admits $\mathcal{O}(\sqrt{T})$ regret \rightarrow FTRL (z_k) admits $\mathcal{O}(1)$ regret

2 action simplex

- If Alice admits 2 actions $\rightarrow \mathcal{O}(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - ▶ Greedy Best Response $\rightarrow w_k = \arg \min_{x \in \Delta_n} [x^\top A y_{2k}]$
 - ▶ Follow the Regularized Leader \rightarrow

$$z_k = \arg \min_{x \in \Delta_n} \left[\sum_{k=1}^{t/2} x^\top A y_{2k} + \gamma \|x\|^2 \right]$$

- ▶ If FTRL (z_k) admits $\mathcal{O}(\sqrt{T})$ regret \rightarrow Greedy BR (w_k) admits $\mathcal{O}(1)$ regret
 - ▶ If Greedy BR (w_k) admits $\mathcal{O}(\sqrt{T})$ regret \rightarrow FTRL (z_k) admits $\mathcal{O}(1)$ regret
- Alice plays x_{2k+1} in the convex hull of z_k and w_k .

Take-Away Message



Algorithms beyond $\mathcal{O}(\sqrt{T})$ regret lower bounds of simultaneous play.

	Simultaneous Play	Alternating Play
General Games	$\tilde{\mathcal{O}}(\sqrt{T})$ tight	$\tilde{\mathcal{O}}(T^{1/3})$
2 actions for Alice	$\tilde{\mathcal{O}}(\sqrt{T})$ tight	$\tilde{\mathcal{O}}(\log T)$

Take-Away Message



Algorithms beyond $\mathcal{O}(\sqrt{T})$ regret lower bounds of simultaneous play.

	Simultaneous Play	Alternating Play
General Games	$\tilde{\mathcal{O}}(\sqrt{T})$ tight	$\tilde{\mathcal{O}}(T^{1/3})$
2 actions for Alice	$\tilde{\mathcal{O}}(\sqrt{T})$ tight	$\tilde{\mathcal{O}}(\log T)$

Thank you!!

References
