

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

Payo matrix of Alice

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

Payo 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

Payo matrix A of Alice

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}

Payo matrix B of Bob

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

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.

$$x = (x_1, \dots, x_n) \quad n$$

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.

$$y = (y_1, \dots, y_m) \quad m$$

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.

$$x = (x_1, \dots, x_n) \quad n$$

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}

Alice's expected cost

$$x \cdot Ay$$

Bob plays a prob. distr.

$$y = (y_1, \dots, y_m) \quad m$$

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}

Bob's expected cost

$$x \cdot By$$

Normal Form Games

Normal-Form Games - Mixed Strategies

Matrices A and B

Alice plays a prob. distr.

$$x = (x_1, \dots, x_n) \quad n$$

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}

$$x^T A y$$

Bob plays a prob. distr.

$$y = (y_1, \dots, y_m) \quad m$$

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}

$$x^T B y$$

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 – Alice gets $x_1 A y_1$

Simultaneous Play

Simultaneous Play

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

- Alice selects x_1 / Bob selects y_1 – Alice gets $x_1 A y_1$
- 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 – Alice gets $x_1 A y_1$
- Alice selects x_2 / Bob selects y_2 – Alice gets $x_2 A y_2$

Simultaneous Play

Simultaneous Play

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

Alice selects x_1 / **Bob** selects y_1 ! **Alice** gets $x_1^> A y_1$

Alice selects x_2 / **Bob** selects y_2 ! **Alice** gets $x_2^> A y_2$

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 ! **Alice** gets $x_1^> A y_1$

Alice selects x_2 / **Bob** selects y_2 ! **Alice** gets $x_2^> A y_2$

Alice selects x_3 / **Bob** selects y_3 ! **Alice** gets $x_3^> A y_3$

Simultaneous Play

Simultaneous Play

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

Alice selects x_1 / **Bob** selects y_1 ! **Alice** gets $x_1^T A y_1$

Alice selects x_2 / **Bob** selects y_2 ! **Alice** gets $x_2^T A y_2$

Alice selects x_3 / **Bob** selects y_3 ! **Alice** gets $x_3^T A y_3$

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 ! **Alice** gets $x_1^T A y_1$

Alice selects x_2 / **Bob** selects y_2 ! **Alice** gets $x_2^T A y_2$

Alice selects x_3 / **Bob** selects y_3 ! **Alice** gets $x_3^T A y_3$

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 ! Alice gets $x_1^T A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^T A y_2$

Simultaneous play

Simultaneous Play

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

Alice selects x_1 / Bob selects y_1 ! Alice gets $x_1^T y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^T y_2$

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

$$x_{t+1}(i) = \frac{x_t(i) e^{[A y_t]_i}}{\sum_{j=1}^n x_t(j) e^{[A y_t]_j}}$$

Simultaneous play

Simultaneous Play

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

Alice selects x_1 / Bob selects y_1 ! Alice gets $x_1^T A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^T A y_2$

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

$$x_{t+1}(i) = \frac{e^{[A y_t]_i} x_t(i)}{\sum_{j=1}^n e^{[A y_t]_j} x_t(j)}$$

Theorem (Freund and Schapire JCSS '97)

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

$$R(T) := \sum_{t=1}^T \{ \text{cost of Alice} \} - \min_{i \in [n]} \sum_{t=1}^T \{ \text{cost of best action} \} \leq \sqrt{T}$$

Simultaneous play

Simultaneous Play

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

Alice selects x_1 / Bob selects y_1 ! Alice gets $x_1^T A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^T A y_2$

Alice selects x_3 / Bob selects y_3 ! Alice gets $x_3^T A y_3$

Theorem (Freund and Schapire JCSS '97)

No matter Bob's strategies,

$$\frac{1}{T} \sum_{t=1}^T x_t^T A y_t$$

time-average cost

$$\frac{1}{T} \min_{x \in \mathcal{X}} \sum_{t=1}^T x^T A y_t$$

best fixed action

$$+ O\left(\frac{\sqrt{pT}}{T}\right) \rightarrow 0$$

Simultaneous play

Simultaneous Play

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

Alice selects x_1 / Bob selects y_1 ! Alice gets $x_1^T A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^T A y_2$

Theorem (Freund and Schapire JCSS '97)

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

$$R(T) := \sum_{t=1}^T x_t^T A y_t - \min_{i \in [n]} \sum_{t=1}^T [A y_t]_i \leq \sigma \sqrt{T}$$

cost of Alice
cost of best action

Simultaneous play

Simultaneous Play

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

Alice selects x_1 / Bob selects y_1 ! Alice gets $x_1^T A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^T A y_2$

Theorem (Freund and Schapire JCSS '97)

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

$$R(T) := \sum_{t=1}^T x_t^T A y_t - \min_{z \in \{z\}} \sum_{t=1}^T z^T A y_t = O(\sqrt{T})$$

cost of Alice
cost of best action

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 ! Alice gets $x_1^T A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^T A y_2$

Theorem (Freund and Schapire JCSS '97)

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

$$R(T) := \sum_{t=1}^T x_t^T A y_t - \min_{z \in \{z\}} \sum_{t=1}^T z^T A y_t \leq \sigma \sqrt{T}$$

cost of Alice
cost of best action

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 ! Alice gets $x_1^\top A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^\top A y_2$

Theorem (folklore)

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

$$R(T) := \sum_{t=1}^T x_t^\top A y_t - \min_{i \in [n]} \sum_{t=1}^T [A y_t]_i \quad (O(\sqrt{T}))$$

$\left\{ \begin{array}{l} \text{cost of Alice} \end{array} \right\}$

 $\left\{ \begin{array}{l} \text{cost of best action} \end{array} \right\}$

$O(\sqrt{T})$ 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 ! Alice gets $x_1^\top A y_1$

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^\top A y_2$

Theorem (folklore)

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

$$R(T) := \sum_{t=1}^T x_t^\top A y_t - \min_{\substack{i \in \{1, \dots, n\} \\ z \in \{z\}}} \sum_{t=1}^T [A y_t]_i \quad (p \bar{T})$$

cost of Alice
cost of best action

Simultaneous play

Simultaneous Play

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

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

Alice selects x_2 / Bob selects y_2 ! Alice gets $x_2^\top A y_2$

Theorem (folklore)

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

$$R(T) := \sum_{t=1}^T x_t^\top A y_t - \min_{\substack{i \in [n] \\ \{z\}}} \sum_{t=1}^T [A y_t]_i \quad (p \bar{T})$$

cost of Alice
cost of best action

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 ! **Alice** gets $x_1^{\geq} A y_0$

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 ! Alice gets $x_1^T A y_0$

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 ! **Alice** gets $x_1^> A y_0$

Bob selects y_2 ! **Alice** gets $x_1^> A y_2$

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 ! **Alice** gets $x_1^> A y_0$

Bob selects y_2 ! **Alice** gets $x_1^> A y_2$

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 ! **Alice** gets $x_1^> A y_0$

Bob selects y_2 ! **Alice** gets $x_1^> A y_2$

Alice selects x_3 ! **Alice** gets $x_3^> A y_2$

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 ! **Alice** gets $x_1^> A y_0$

Bob selects y_2 ! **Alice** gets $x_1^> A y_2$

Alice selects x_3 ! **Alice** gets $x_3^> A y_2$

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 ! **Alice** gets $x_1^> A y_0$

Bob selects y_2 ! **Alice** gets $x_1^> A y_2$

Alice selects x_3 ! **Alice** gets $x_3^> A y_2$

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

$$R(T) := \max_{k=0}^{T-2} \min_{x^>} \left\{ \sum_{k=0}^{T-2} A(y_{2k} + y_{2k+2}) \right\} - \min_{k=0}^{T-2} \left\{ \sum_{k=0}^{T-2} A(y_{2k} + y_{2k+2}) \right\}$$

Alice's cost
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 ! Alice gets $x_1^T A y_0$

Bob selects y_2 ! Alice gets $x_1^T A y_2$

Alice selects x_3 ! Alice gets $x_3^T A y_2$

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

In alternating play , Alice can always guarantee

$\mathcal{O}(T^{1/3})$ 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 ! $O(1)$ regret [Bailey et al. COLT 2020]

If both Alice and Bob use Hedge in zero-sum games ! $O(T^{1/3})$ regret [Wibisono et al. NeurIPS 2022]

Alternating Turns

Previous Results

If both Alice and Bob use Gradient Descent in unconstrained zero-sum games ! $O(1)$ regret [Bailey et al. COLT 2020]

If both Alice and Bob use Hedge in zero-sum games ! $O(T^{1/3})$ 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 ! $O(\sqrt{T})$ regret [Bailey et al. COLT 2020]

If both Alice and Bob use Hedge in zero-sum games ! $O(\sqrt{T})$ regret [Wibisono et al. NeurIPS 2022]

Previous results make assumptions on Bob's behavior.

Theorem ([Cevher, Cutksovsky[?], Piliouras, Skoulakis[?], Viano NeurIPS '23 spotlight])

In alternating play, Alice can always guarantee

$O(\sqrt{T})$ regret for general games.

$O(\log T)$ 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 ! **Alice** gets $x_1^> A y_0$

Bob selects y_2 ! **Alice** gets $x_1^> A y_2$

Alice selects x_3 ! **Alice** gets $x_3^> A y_2$

	Simultaneous Play	Alternating Play
General Games	$\Theta\left(\binom{P}{T}\right)$ tight	$\Theta(T^{1=3})$
2 actions for Alice	$\Theta\left(\binom{P}{T}\right)$ tight	$\Theta(\log T)$

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 ! Alice gets $x_1^T A y_0$

Bob selects y_2 ! Alice gets $x_1^T A y_2$

Alice selects x_3 ! Alice gets $x_3^T A y_2$

Our algorithm (general simplex)

$$x_{2k+1} = \underset{x \in \Delta_n}{\operatorname{argmin}} \left[\sum_{i=1}^n x_i \log x_i + \sum_{k=0}^k x^T A y_{2k} + \sum_{i=1}^n x_i \log x_i \right]$$

exploits y_{2k}
reinforces good past actions
prevents over fitting

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 – Alice gets $x_1 A y_0$
- Bob selects y_2 – Alice gets $x_1 A y_2$
- Alice selects x_3 – Alice gets $x_3 A y_2$
- ...

Our algorithm (general simplex)

$$x_{2k+1} = \underset{x}{\operatorname{argmin}} \left[\sum_{i=1}^n x_i A y_{2k} + 2 \sum_{k=0}^{k-1} x_i A y_{2k} - \sum_{i=1}^n \log x_i \right]$$

exploits y_{2k}
reinforces good past actions
prevents overfitting

Setting $\epsilon = O(T^{-1/3})$ $O(T^{1/3})$ alternating regret

2 action simplex

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

2 action simplex

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

2 action simplex

- If Alice admits 2 actions $O(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - ┆ Greedy Best Response $w_k = \arg \min_x \sum_n x \cdot Ay_{2k}$

2 action simplex

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

$$z_k = \arg \min_{x \in \Delta_n} \sum_{k=1}^{t/2} x \cdot Ay_{2k} + \|x\|^2$$

2 action simplex

- If Alice admits 2 actions $O(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - | Greedy Best Response $w_k = \arg \min_{x \in \mathcal{X}} \sum_{n=1}^k x \cdot Ay_{2k}$
 - | Follow the Regularized Leader

$$z_k = \arg \min_{x \in \mathcal{X}} \sum_{n=1}^{t/2} x \cdot Ay_{2k} + \frac{1}{2} \|x\|^2$$

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

2 action simplex

- If Alice admits 2 actions $O(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - | Greedy Best Response $w_k = \arg \min_{x \in \mathcal{X}} \sum_{n=1}^k x \cdot A y_{2k}$
 - | Follow the Regularized Leader

$$z_k = \arg \min_{x \in \mathcal{X}} \sum_{n=1}^{t/2} x \cdot A y_{2k} + \frac{1}{2} \|x\|^2$$

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

2 action simplex

- If Alice admits 2 actions $O(\log T)$ in alternating play!
- At round $t = 2k + 1$
 - | Greedy Best Response $w_k = \arg \min_{x \in \Delta_n} \sum_{i=1}^n x_i A_{i,2k}$
 - | Follow the Regularized Leader

$$z_k = \arg \min_{x \in \Delta_n} \sum_{i=1}^n x_i A_{i,2k} + \frac{t}{2} \|x\|^2$$

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

Take-Away Message

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

	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	$\tilde{O}(\log T)$

Take-Away Message

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

	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	$\tilde{O}(\log T)$

Thank you!!

References
