Learning in Games with Conflicting Interests

Davide Legacci

Université Grenoble Alpes

October 30, 2024 – CJC-MA, Lyon

About

- L, Mertikopoulos, & Pradelski, **A Geometric Decomposition of Finite Games: Convergence vs. Recurrence under Exponential Weights**, ICML2024
- L, Mertikopoulos, Papadimitriou, Piliouras, & Pradelski, **No-regret Learning in Harmonic Games: Extrapolation in the Face of Conflicting Interests**, NeurIPS2024

   

This talk explained to my dad

- **Game**: **Strategic** interaction among **players** aiming at maximizing their **utility**.
- **Equilibrium**: No player can improve their utility by changing their own strategy.
  - *Common interest*: There exists an equilibrium all players are interested in reaching.
  - *Conflicting interests*: Players can't find an agreement. *"No matter what you do – I want the opposite"*.
- **Learning in games**: Do players learn to **reach an equilibrium** through **repeated interaction**?
  - *Common interest*: **Yes!** Rational players eventually learn to adopt an equilibrium strategy.
  - *Conflicting interests*: *Not much was known...*

Theorem (Informal, Legacci et al., 2024)

Standard **learning** *procedures in continuous time are* **quasi-periodic** *in games with* **conflicting interests**.

This talk explained to my dad

- **Game**: **Strategic** interaction among **players** aiming at maximizing their **utility**.
- **Equilibrium**: No player can improve their utility by changing their own strategy.
  - *Common interest*: There exists an equilibrium all players are interested in reaching.
  - *Conflicting interests*: Players can't find an agreement. *"No matter what you do – I want the opposite"*.
- **Learning in games**: Do players learn to **reach an equilibrium** through **repeated interaction**?
  - *Common interest*: **Yes!** Rational players eventually learn to adopt an equilibrium strategy.
  - *Conflicting interests*: *Not much was known...*

Theorem (Informal, Legacci et al., 2024)

Standard **learning** *procedures in continuous time are* **quasi-periodic** *in games with* **conflicting interests**.

This talk explained to my dad

- **Game**: **Strategic** interaction among **players** aiming at maximizing their **utility**.
- **Equilibrium**: No player can improve their utility by changing their own strategy.
  - *Common interest*: There exists an equilibrium all players are interested in reaching.
  - *Conflicting interests*: Players can't find an agreement. *"No matter what you do – I want the opposite"*.
- **Learning in games**: Do players learn to **reach an equilibrium** through **repeated interaction**?
  - *Common interest*: **Yes!** Rational players eventually learn to adopt an equilibrium strategy.
  - *Conflicting interests*: *Not much was known...*

Theorem (Informal, Legacci et al., 2024)

Standard **learning** *procedures in continuous time are* **quasi-periodic** *in games with* **conflicting interests**.

This talk explained to my dad

- **Game**: **Strategic** interaction among **players** aiming at maximizing their **utility**.
- **Equilibrium**: No player can improve their utility by changing their own strategy.
  - *Common interest*: There exists an equilibrium all players are interested in reaching.
  - *Conflicting interests*: Players can't find an agreement. *"No matter what you do – I want the opposite"*.
- **Learning in games**: Do players learn to **reach an equilibrium** through **repeated interaction**?
  - *Common interest*: **Yes!** Rational players eventually learn to adopt an equilibrium strategy.
  - *Conflicting interests*: *Not much was known...*

Theorem (Informal, Legacci et al., 2024)

Standard **learning** *procedures in continuous time are* **quasi-periodic** *in games with* **conflicting interests**.

This talk explained to my dad

- **Game**: **Strategic** interaction among **players** aiming at maximizing their **utility**.
- **Equilibrium**: No player can improve their utility by changing their own strategy.
  - *Common interest*: There exists an equilibrium all players are interested in reaching.
  - *Conflicting interests*: Players can't find an agreement. *"No matter what you do – I want the opposite"*.
- **Learning in games**: Do players learn to **reach an equilibrium** through **repeated interaction**?
  - *Common interest*: **Yes!** Rational players eventually learn to adopt an equilibrium strategy.
  - *Conflicting interests*: *Not much was known...*

Theorem (Informal, Legacci et al., 2024)

Standard **learning** *procedures in continuous time are* **quasi-periodic** *in games with* **conflicting interests**.

Question (1.0)

What is the long-run behavior of **learning dynamics** *in games where players have* **conflicting interests***?*

1. **Games with conflicting interests**

2. **Learning in games**

3. **Learning in games with conflicting interests**

1. Games with conflicting interests

2. Learning in games

3. Learning in games with conflicting interests

Taxonomy of Games

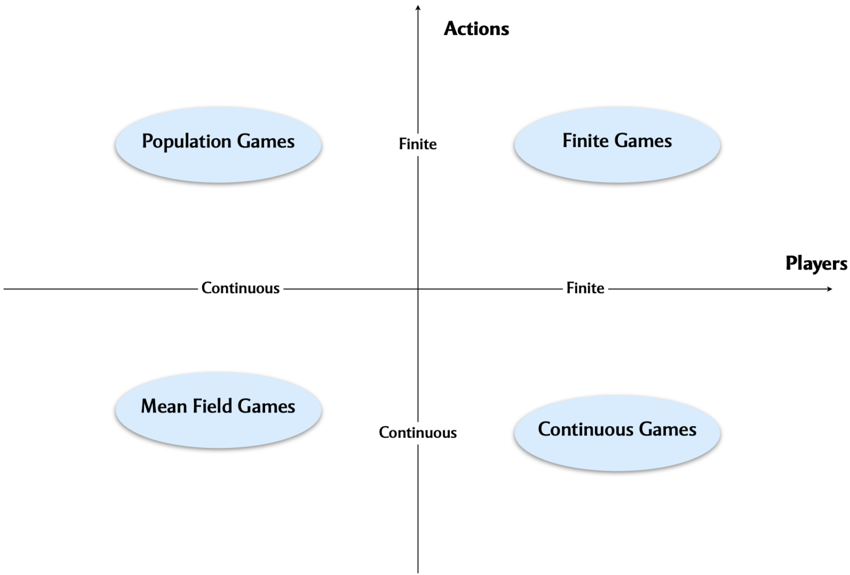

Taxonomy of Games

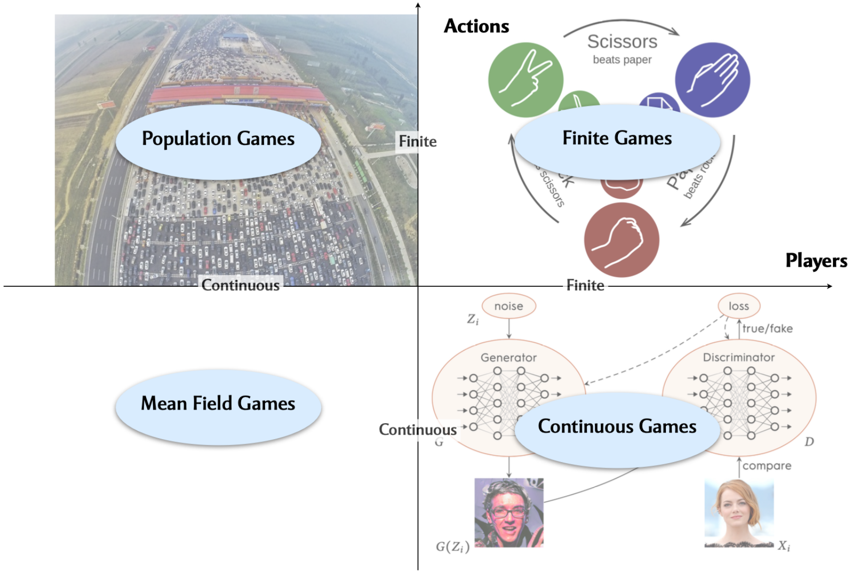

Finite Games and their Mixed Extension

Definition

A **finite game** $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- A finite set of *players* $\mathcal{N} = \{1, \dots, N\}$
- A finite set of *actions* $\mathcal{A}_i = \{1, \dots, A_i\}$ per player $i \in \mathcal{N}$
- An ensemble of *payoff functions* $u_i : \mathcal{A} \equiv \prod_{j \in \mathcal{N}} \mathcal{A}_j \to \mathbb{R}$, $i \in \mathcal{N}$ # *Each player's payoff depends on everybody's actions.*

Definition

The **mixed extension** of a finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- The set of *mixed strategies* $x_i \in \mathcal{X}_i := \Delta(\mathcal{A}_i)$ per player $i \in \mathcal{N}$, discrete prob. dist. over pure strategies
- The ensemble of *expected payoffs* $u_i : \mathcal{X} \equiv \prod_{j \in \mathcal{N}} \mathcal{X}_j \to \mathbb{R}$ defined by $u_i(x) = \mathbb{E}_{\alpha \sim x}[u_i(\alpha)]$

Definition

The **payoff field** is the array of differentials of the payoff functions *wrt. the variables of the respective player*:

$$v_i(x) = d_i u_i(x) \quad \text{for all } i \in \mathcal{N}, x \in \mathcal{X} \quad \rightsquigarrow \quad \text{1-form, generically } \textbf{not exact}$$

Finite Games and their Mixed Extension

Definition

A **finite game** $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- A finite set of *players* $\mathcal{N} = \{1, \dots, N\}$
- A finite set of *actions* $\mathcal{A}_i = \{1, \dots, A_i\}$ per player $i \in \mathcal{N}$
- An ensemble of *payoff functions* $u_i : \mathcal{A} \equiv \prod_{j \in \mathcal{N}} \mathcal{A}_j \to \mathbb{R}$, $i \in \mathcal{N}$ # *Each player's payoff depends on everybody's actions.*

Definition

The **mixed extension** of a finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- The set of *mixed strategies* $x_i \in \mathcal{X}_i := \Delta(\mathcal{A}_i)$ per player $i \in \mathcal{N}$, discrete prob. dist. over pure strategies
- The ensemble of *expected payoffs* $u_i : \mathcal{X} \equiv \prod_{j \in \mathcal{N}} \mathcal{X}_j \to \mathbb{R}$ defined by $u_i(x) = \mathbb{E}_{\alpha \sim x}[u_i(\alpha)]$

Definition

The **payoff field** is the array of differentials of the payoff functions *wrt. the variables of the respective player*:

$$v_i(x) = d_i u_i(x) \quad \text{for all } i \in \mathcal{N}, x \in \mathcal{X} \quad \rightsquigarrow \quad \text{1-form, generically } \textbf{not exact}$$

Finite Games and their Mixed Extension

Definition

A **finite game** $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- A finite set of *players* $\mathcal{N} = \{1, \dots, N\}$
- A finite set of *actions* $\mathcal{A}_i = \{1, \dots, A_i\}$ per player $i \in \mathcal{N}$
- An ensemble of *payoff functions* $u_i : \mathcal{A} \equiv \prod_{j \in \mathcal{N}} \mathcal{A}_j \to \mathbb{R}$, $i \in \mathcal{N}$ # *Each player's payoff depends on everybody's actions.*

Definition

The **mixed extension** of a finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- The set of *mixed strategies* $x_i \in \mathcal{X}_i := \Delta(\mathcal{A}_i)$ per player $i \in \mathcal{N}$, discrete prob. dist. over pure strategies
- The ensemble of *expected payoffs* $u_i : \mathcal{X} \equiv \prod_{j \in \mathcal{N}} \mathcal{X}_j \to \mathbb{R}$ defined by $u_i(x) = \mathbb{E}_{\alpha \sim x}[u_i(\alpha)]$

Definition

The **payoff field** is the array of differentials of the payoff functions *wrt. the variables of the respective player*:

$$v_i(x) = d_i u_i(x) \quad \text{for all } i \in \mathcal{N}, x \in \mathcal{X} \quad \rightsquigarrow \quad \text{1-form, generically } \textbf{not exact}$$

Finite Games and their Mixed Extension

Definition

A finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- A finite set of players $\mathcal{N} = \{1, \dots, N\}$
- A finite set of actions $\mathcal{A}_i = \{1, \dots, A_i\}$ per player $i \in \mathcal{N}$
- An ensemble of payoff functions $u_i : \mathcal{A} \equiv \prod_{j \in \mathcal{N}} \mathcal{A}_j \to \mathbb{R}$, $i \in \mathcal{N}$

Definition

The mixed extension of a finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ consists of

- The set of mixed strategies $x_i \in \mathcal{X}_i := \Delta(\mathcal{A}_i)$ per player $i \in \mathcal{N}$, discrete prob. dist. over pure strategies
- The ensemble of expected payoffs $u_i : \mathcal{X} \equiv \prod_{j \in \mathcal{N}} \mathcal{X}_j \to \mathbb{R}$ defined by $u_i(x) = \mathbb{E}_{\alpha \sim x}[u_i(\alpha)]$

Definition

The **payoff field** is the array of differentials of the payoff functions wrt. the variables of the respective player:

$$v_i(x) = d_i u_i(x) \quad \text{for all } i \in \mathcal{N}, x \in \mathcal{X} \quad \rightsquigarrow \quad \text{1-form, generically } \textbf{not exact}$$

Games with Common & Conflicting Interests

Definition (Potential games [MS96; San10])

Finite game with **exact** payoff field: $v = d\phi$.

Games with Common & Conflicting Interests

Definition (Potential games [MS96; San10])

Finite game with **exact** payoff field: $v = d\phi$.

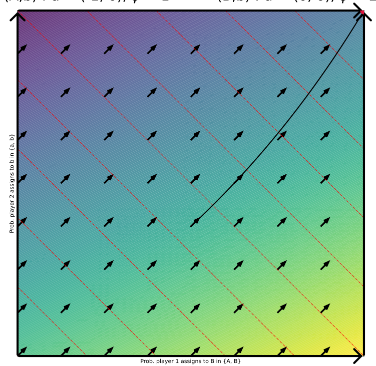

Games with Common & Conflicting Interests

Definition (Potential games [MS96; San10])

Finite game with **exact** payoff field: $v = d\phi$.

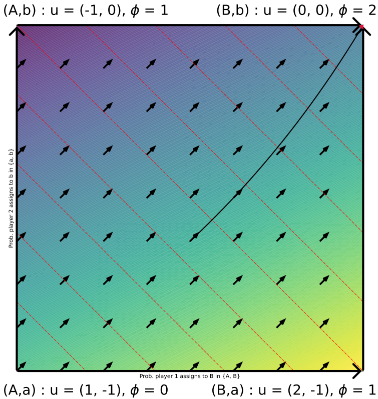

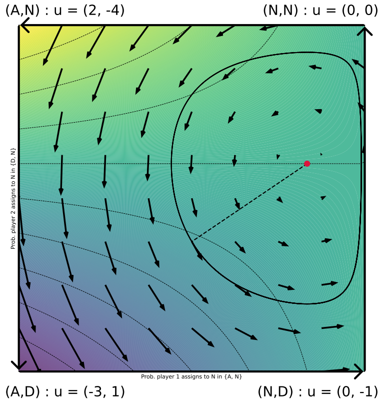

Games with Common & Conflicting Interests

Definition (Potential games [MS96; San10])

Finite game with **exact** payoff field: $v = d\phi$.

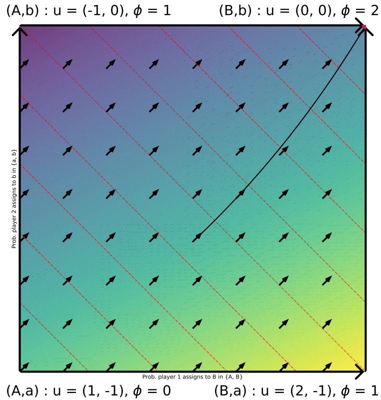

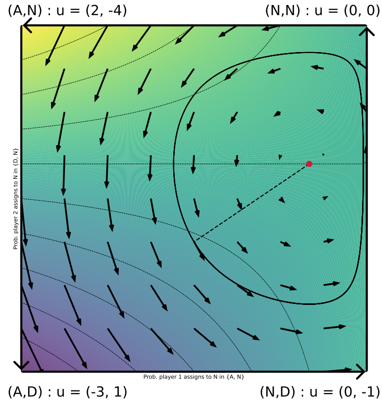

Definition (Harmonic games [Can+11; Abd+22; Leg+24])

Finite game with $m \in \mathbb{R}^N_{++}$, $x^* \in \operatorname{rint} \mathcal{X}$ such that $\langle v(x), x - x^* \rangle_m = 0$ for all $x \in \mathcal{X}$.

Harmonic games to model conflicting interests

Framework: combinatorial Hodge theory.

- Two players: an **army** and a **fortress**. The army can **Attack** or **Not**; the fortress can **Defend** or **Not**.
- Tacking an action (A or D) implies a **preparation cost**.
- The army prevails attacking the undefended fortress; however, if the fortress is defended, the assault fails.

Harmonic games to model conflicting interests

- Two players: an **army** and a **fortress**. The army can **Attack** or **Not**; the fortress can **Defend** or **Not**.
- Tacking an action (A or D) implies a **preparation cost**.
- The army prevails attacking the undefended fortress; however, if the fortress is defended, the assault fails.

	D	N
A	$-3, 1$	$2, -4$
N	$0, -1$	$0, 0$

Harmonic games to model conflicting interests

- Two players: an **army** and a **fortress**. The army can **Attack** or **Not**; the fortress can **Defend** or **Not**.
- Tacking an action (A or D) implies a **preparation cost**.
- The army prevails attacking the undefended fortress; however, if the fortress is defended, the assault fails.
- **Not a zero-sum game!**

	D	N
A	$-3, 1$	$2, -4$
N	$0, -1$	$0, 0$

Harmonic games to model conflicting interests

- Two players: an **army** and a **fortress**. The army can **Attack** or **Not**; the fortress can **Defend** or **Not**.
- Tacking an action (A or D) implies a **preparation cost**.
- The army prevails attacking the undefended fortress; however, if the fortress is defended, the assault fails.
- **Not a zero-sum game!**

	D	N
A	$-3, 1$	$2, -4$
N	$0, -1$	$0, 0$

Harmonic game: $\sum_{i \in \mathcal{N}} m_i \langle v_i(x), x_i - x_i^* \rangle = 0$ for all $x \in \mathcal{X}$ with $m = (6, 5)$ and $x^* = \left(\left(\frac{1}{6}, \frac{5}{6}\right), \left(\frac{2}{5}, \frac{3}{5}\right) \right)$.

Harmonic games to model conflicting interests

- Two players: an **army** and a **fortress**. The army can **Attack** or **Not**; the fortress can **Defend** or **Not**.
- Tacking an action (A or D) implies a **preparation cost**.
- The army prevails attacking the undefended fortress; however, if the fortress is defended, the assault fails.
- **Not a zero-sum game!**

	D	N
A	$-3, 1$	$2, -4$
N	$0, -1$	$0, 0$

Harmonic game: $\sum_{i \in \mathcal{N}} m_i \langle v_i(x), x_i - x_i^* \rangle = 0$ for all $x \in \mathcal{X}$ with $m = (6, 5)$ and $x^* = \left(\left(\frac{1}{6}, \frac{5}{6}\right), \left(\frac{2}{5}, \frac{3}{5}\right)\right)$.

Example (Other harmonic games)

- Matching Pennies, Rock-Paper-Scissors
- All two-player zero-sum games with an interior equilibrium — Legacci et al. (2024)
- Cyclic games — Hofbauer and Schlag (2000)

Harmonic games to model conflicting interests

- Two players: an **army** and a **fortress**. The army can **Attack** or **Not**; the fortress can **Defend** or **Not**.
- Tacking an action (A or D) implies a **preparation cost**.
- The army prevails attacking the undefended fortress; however, if the fortress is defended, the assault fails.
- **Not a zero-sum game!**

	D	N
A	$-3, 1$	$2, -4$
N	$0, -1$	$0, 0$

Harmonic game: $\sum_{i \in \mathcal{N}} m_i \langle v_i(x), x_i - x_i^* \rangle = 0$ for all $x \in \mathcal{X}$ with $m = (6, 5)$ and $x^* = ((\frac{1}{6}, \frac{5}{6}), (\frac{2}{5}, \frac{3}{5}))$.

Theorem (Decomposition [Can+11; Abd+22])

The vector space of finite games $\mathcal{U}$ modulo strategic equivalence admits the orthogonal direct sum decomposition

$$\mathcal{U} = \mathcal{U}_{potential} \oplus \mathcal{U}_{harmonic}$$

Why we care about harmonic games

- Natural **complement to potential games** from a strategic viewpoint
- Benchmark for strategic interaction with **conflicting interests**, generalizing zero-sum games to N-player setting

How does the "circular" payoff structure of games in which players have conflicting interests affect learning?

Question (2.0)

What is the long-run behavior of **learning dynamics** *in* **harmonic games***?*

Why we care about harmonic games

- Natural **complement to potential games** from a strategic viewpoint
- Benchmark for strategic interaction with **conflicting interests**, generalizing zero-sum games to N-player setting

How does the "circular" payoff structure of games in which players have conflicting interests affect learning?

Question (2.0)

What is the long-run behavior of **learning dynamics** *in* **harmonic games***?*

Multi-agent learning

Do players learn to emulate rational behaviour through repeated interactions?

Multi-objective optimization

- each agent aims at maximizing their payoff
- each agent's payoff depends on everybody's actions

Multi-agent learning

Multi-agent learning;

for *each epoch and every player ;*

do

Choose *action* ;

Receive *reward* ;

Get *feedback* (maybe) ;

Defining elements

- **Time:** continuous or discrete?
- **Players:** continuous or finite?
- **Actions:** continuous or finite? constrained or unconstrained?
- **Rewards:** endogenous or exogenous (determined by other players or by "Nature")?
- **Feedback:** full information, incomplete (bandit) information? Stochastic noise?

Multi-agent learning

Learning in finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$;

for *each* $t \geq 0$ *and every player* $i \in \mathcal{N}$;

do simultaneously

Choose *mixed strategy* $x_i(t) \in \mathcal{X}_i = \Delta(\mathcal{A}_i)$;

Receive *mixed payoff* $u_i(x(t)) \in \mathbb{R}$;

Get *payoff vector* $v_i(x(t)) = d_i u_i(x(t)) \in \mathbb{R}^{A_i}$;

Defining elements

- **Time:** continuous
- **Players:** finite
- **Actions:** continuous, constrained
- **Rewards:** endogenous
- **Feedback:** full information

Multi-agent learning

Learning in finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$;

for *each* $t \geq 0$ *and every player* $i \in \mathcal{N}$ *;*

do simultaneously

Choose *mixed strategy* $x_i(t) \in \mathcal{X}_i = \Delta(\mathcal{A}_i)$;

Receive *mixed payoff* $u_i(x(t)) \in \mathbb{R}$;

Get *payoff vector* $v_i(x(t)) = d_i u_i(x(t)) \in \mathbb{R}^{A_i}$;

Defining elements

- **Time:** continuous
- **Players:** finite
- **Actions:** continuous, constrained
- **Rewards:** endogenous
- **Feedback:** full information

Key: choice of next strategy given previous payoff vector.

Regularized learning & Hessian-Riemannian interpretation

Key: choice of next strategy given previous payoff vector.

- Score each action α_i based on its cumulative effect on payoff over time
- Play an action with probability given by soft argmax

Follow-the-regularized-leader (FTRL) — **Shalev-shwartz and Singer (2006)**

$$\dot{y}_i(t) = v_i(x(t)), \qquad x_i(t) = \nabla h_i^*(y_i(t)) = \arg\max_{z_i \in \mathcal{X}_i} \{\langle y_i(t), z_i \rangle - h_i(z_i)\}$$

where $h_i : \mathcal{X}_i \to \mathbb{R}$ is strongly convex regularizer.

Regularized learning & Hessian-Riemannian interpretation

Key: choice of next strategy given previous payoff vector.

- Score each action α_i based on its cumulative effect on payoff over time
- Play an action with probability given by soft argmax

Follow-the-regularized-leader (FTRL) **Shalev-shwartz and Singer (2006)**

$$\dot{y}_i(t) = v_i(x(t)), \qquad x_i(t) = \nabla h_i^*(y_i(t)) = \arg\max_{z_i \in \mathcal{X}_i} \{\langle y_i(t), z_i \rangle - h_i(z_i)\}$$

where $h_i : \mathcal{X}_i \to \mathbb{R}$ is strongly convex regularizer.

Hessian-Riemannian individual gradient ascent **Mertikopoulos and Sandholm (2018)**

$$\dot{x}_i = v_i^{\sharp}(x) = \operatorname{grad}_i u_i(x) \quad \text{with respect to Hessian metric } g = \operatorname{Hess} h \text{ on } \operatorname{rint} \mathcal{X}$$

$\rightsquigarrow$ *Every player follows direction of maximal payoff increase*

Regularizer of Legendre type

Regularized learning & Hessian-Riemannian interpretation

Follow-the-regularized-leader (FTRL) — Shalev-shwartz and Singer (2006)

$$\dot{y}_i(t) = v_i(x(t)), \qquad x_i(t) = \nabla h_i^*(y_i(t)) = \arg\max_{z_i \in \mathcal{X}_i} \{\langle y_i(t), z_i \rangle - h_i(z_i)\}$$

where $h_i : \mathcal{X}_i \to \mathbb{R}$ is strongly convex regularizer.

Hessian-Riemannian individual gradient ascent — Mertikopoulos and Sandholm (2018)

$$\dot{x}_i = {v_i}^\sharp(x) = \operatorname{grad}_i u_i(x) \quad \text{with respect to Hessian metric } g = \operatorname{Hess} h \text{ on } \operatorname{rint} \mathcal{X}$$

⇝ *Every player follows direction of maximal payoff increase*

Regularizer of Legendre type

Example (Canonical: Euclidean projection dynamics [Fri91])

$$\text{Euclidean regularizer } h = \frac{1}{2}\|x\|_2^2 \quad \Longrightarrow \quad x_i(t) = \arg\min_{z_i \in \mathcal{X}_i} \|y_i(t) - z_i\|_2^2$$

Regularized learning & Hessian-Riemannian interpretation

Follow-the-regularized-leader (FTRL) — **Shalev-shwartz and Singer (2006)**

$$\dot{y}_i(t) = v_i(x(t)), \qquad x_i(t) = \nabla h_i^*(y_i(t)) = \arg\max_{z_i \in \mathcal{X}_i} \{\langle y_i(t), z_i \rangle - h_i(z_i)\}$$

where $h_i : \mathcal{X}_i \to \mathbb{R}$ is strongly convex regularizer.

Hessian-Riemannian individual gradient ascent — **Mertikopoulos and Sandholm (2018)**

$$\dot{x}_i = v_i^\sharp(x) = \operatorname{grad}_i u_i(x) \quad \text{with respect to Hessian metric } g = \operatorname{Hess} h \text{ on } \operatorname{rint} \mathcal{X}$$

$\rightsquigarrow$ *Every player follows direction of maximal payoff increase*

Regularizer of Legendre type

Example (Canonical: Exponential weights – Replicator dynamics [Aue+95; TJ78; LW94])

$$\text{Entropic regularizer } h = x \cdot \log x \quad \Longrightarrow \quad x_{i\alpha_i}(t) = \frac{\exp y_{i\alpha_i}}{\sum_{\beta_i} \exp y_{i\beta_i}} \quad \Longrightarrow \quad \dot{x}_{i\alpha_i} = x_{i\alpha_i}\left[u_i(\alpha_i, x_{-i}) - u_i(x)\right]$$

Learning in harmonic games – How do the trajectories look like?

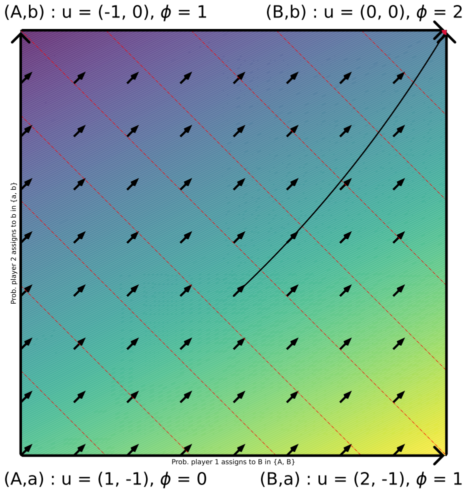

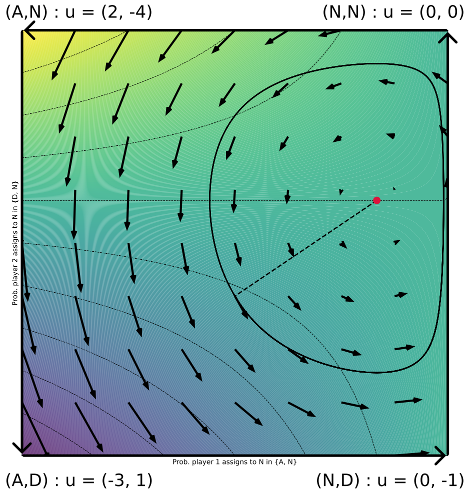

Learning in harmonic games – How do the trajectories look like?

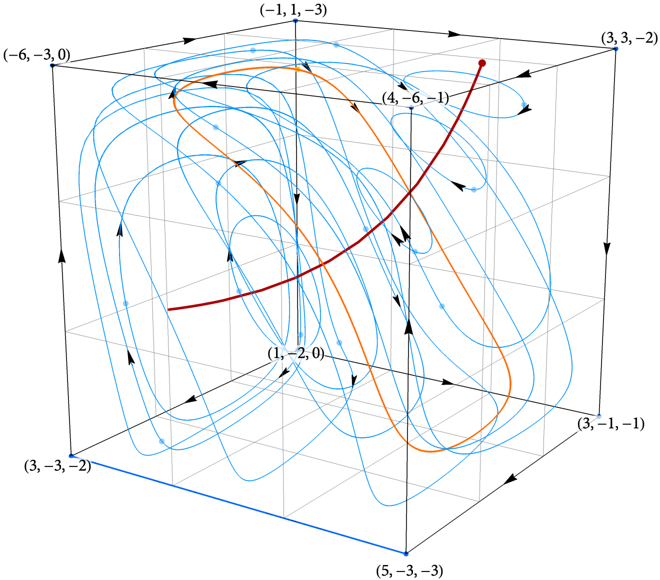

Learning in harmonic games – How do the trajectories look like?

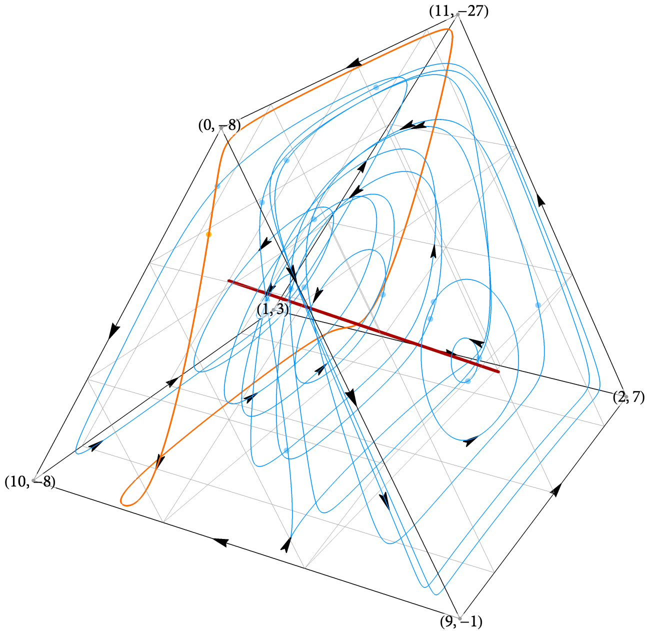

Learning in harmonic games – How do the trajectories look like?

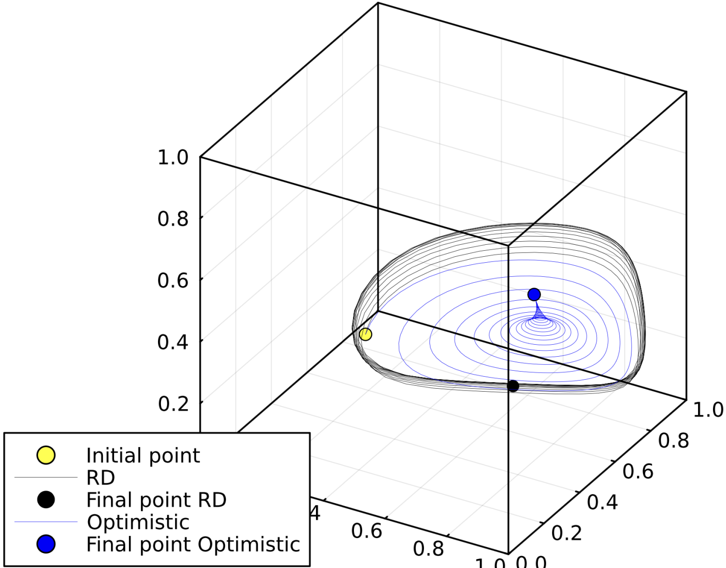

FTRL is globally converging in **potential** games [GP20; HCM17].

Question (3.0)

*What is the long-run behavior of **FTRL** in **harmonic games**?*

Learning in harmonic games – Main result

Definition (Poincaré recurrent dynamical system)

Almost all trajectories return **arbitrarily close** to their starting point **infinitely often**. # *Formalize "quasi-periodicity".*

Theorem (Legacci et al. (2024))

The dynamics of **FTRL** *are* **Poincaré recurrent** *in any* **harmonic game***.*

- Generalize [MPP18] for 2-player zero-sum games with interior equilibrium to *N-player games*

Proof techniques

- **Entropic regularizer:** Replicator dynamics in harmonic games are *volume preserving* under a certain Riemannian metric
- **General FTRL:** Existence of *constant of motion* allows to bound orbits
- In both cases $\rightsquigarrow$ Recurrence by Liouville's and Poincaré's theorems

Learning in harmonic games – Main result

Definition (Poincaré recurrent dynamical system)

Almost all trajectories return **arbitrarily close** to their starting point **infinitely often**. # *Formalize "quasi-periodicity"*.

Theorem (Legacci et al. (2024))

The dynamics of **FTRL** *are* **Poincaré recurrent** *in any* **harmonic game***.*

- Generalize [MPP18] for 2-player zero-sum games with interior equilibrium to *N-player games*

Proof techniques

- **Entropic regularizer:** Replicator dynamics in harmonic games are *volume preserving* under a certain Riemannian metric
- **General FTRL:** Existence of *constant of motion* allows to bound orbits
- In both cases $\rightsquigarrow$ Recurrence by Liouville's and Poincaré's theorems

Learning in harmonic games – Main result

Definition (Poincaré recurrent dynamical system)

Almost all trajectories return **arbitrarily close** to their starting point **infinitely often**. # *Formalize "quasi-periodicity".*

Theorem (Legacci et al. (2024))

The dynamics of **FTRL** *are* **Poincaré recurrent** *in any* **harmonic game**.

- Generalize [MPP18] for 2-player zero-sum games with interior equilibrium to *N-player games*

Proof techniques

- **Entropic regularizer:** Replicator dynamics in harmonic games are *volume preserving* under a certain Riemannian metric
- **General FTRL:** Existence of *constant of motion* allows to bound orbits
- In both cases $\rightsquigarrow$ Recurrence by Liouville's and Poincaré's theorems

Learning in harmonic games – Main result

Definition (Poincaré recurrent dynamical system)

Almost all trajectories return **arbitrarily close** to their starting point **infinitely often**. # *Formalize "quasi-periodicity".*

Theorem (Legacci et al. (2024))

The dynamics of **FTRL** *are* **Poincaré recurrent** *in any* **harmonic game**.

- Generalize [MPP18] for 2-player zero-sum games with interior equilibrium to *N-player games*

Proof techniques

- **Entropic regularizer:** Replicator dynamics in harmonic games are *volume preserving* under a certain Riemannian metric
- **General FTRL:** Existence of *constant of motion* allows to bound orbits
- In both cases $\rightsquigarrow$ Recurrence by Liouville's and Poincaré's theorems

Conclusions

Take-away

Harmonic games complement **potential** games from **strategic** and **dynamic** viewpoints

From here: Learning in discrete time

- **Convergence & good regret guarantees** via optimistic/extra-gradient methods — Legacci et al. (2024)
- Rates of convergence, incomplete/noisy feedback, adaptive step size — *# work in progress*

From here: Learning in continuous time

- Harmonic games with continuous action sets — *# work in progress*
- Stochastic FTRL/RD in harmonic games — *# open direction*

Conclusions

Take-away

Harmonic games complement **potential** games from **strategic** and **dynamic** viewpoints

From here: Learning in discrete time

- **Convergence & good regret guarantees** via optimistic/extra-gradient methods — Legacci et al. (2024)
- Rates of convergence, incomplete/noisy feedback, adaptive step size — *# work in progress*

From here: Learning in continuous time

- Harmonic games with continuous action sets — *# work in progress*
- Stochastic FTRL/RD in harmonic games — *# open direction*

Conclusions

Take-away

Harmonic games complement **potential** games from **strategic** and **dynamic** viewpoints

From here: Learning in discrete time

- **Convergence & good regret guarantees** via optimistic/extra-gradient methods — Legacci et al. (2024)
- Rates of convergence, incomplete/noisy feedback, adaptive step size — # *work in progress*

From here: Learning in continuous time

- Harmonic games with continuous action sets — # *work in progress*
- Stochastic FTRL/RD in harmonic games — # *open direction*

Thanks for your attention!

A glimpse: Harmonic games and Hodge theory

- **Combinatorial** Hodge theory of finite game Γ
- Γ potential iff $\Gamma = d_0\phi \rightsquigarrow$ *exact*
- Γ harmonic iff $\delta_1\Gamma = 0 \rightsquigarrow$ *co-closed*

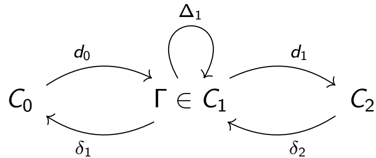

Theorem (Decomposition [Can+11; Abd+22])

$$\Gamma = \Gamma_{potential} + \Gamma_{harmonic}$$

- **Riemannian** setting on rint $\mathcal{X}$; payoff field v of game Γ is 1-form
- Γ potential iff $v = d\phi$
- Riemannian codifferential δ_g : 1-forms (rint $\mathcal{X}$) $\to$ smooth functions (rint $\mathcal{X}$)
- Codifferential: adjoint of differential; dual of divergence.

Theorem (Legacci, Mertikopoulos, and Pradelski (2024))

A finite game Γ is **harmonic iff $\delta_{g^*} v = 0$**, *with respect to certain weights m and a certain metric g^* on* rint $\mathcal{X}$.

Proof sketch - Recurrence of FTRL in harmonic games

Tools: Dynamical systems theory

$$\dot{x} = f(x), \quad f \text{ vector field } : M \text{ open } \subseteq \mathbb{R}^n \to \mathbb{R}$$

- **Liouville's theorem**

$$\operatorname{div} f = 0 \implies \text{volum-preserving system}$$

- **Poincaré's theorem**

$$\begin{cases} \text{volum-preserving system} \\ \text{bounded orbits} \end{cases} \implies \text{recurrent system}$$

Replicator dynamics as Riemannian individual payoff gradient

Recall: payoff field is individual payoff Euclidean gradient

$$v_i(x) = \mathrm{grad}_i\, u_i(x)$$

Replicator field is individual payoff gradient under non-Euclidean geometry g^*: [Sha79]

$$RD_i(x) = \mathrm{grad}_i^{g^*}\, u_i(x)$$

Define divergence operator with respect to geometry g^*

Equivalence between harmonic and divergence-free games

Theorem ([LMP24])

A finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ *is harmonic with uniform measure* $\mu_{i\alpha_i} = 1$ *if and only if its associated replicator vector field* $\operatorname{grad}_i^{g^*} u_i(x)$ *has zero divergence under the geometry* g^**.*

- By Liouville's theorem, RD on harmonic games is volume-preserving in strategy space;
- RD has only bounded orbits in all games;
- Recurrence follows by Poincaré's theorem. □

Riemannian approach: Pros and cons

Pros

- Surprising connection between Riemannian construction and uniform harmonic games
- Fine understanding of dynamics-geometry interplay in strategy space

Cons

- Harmonic / divergence-free equivalence fails changing metric
- Need to change approach for general FTRL case.

For general FTRL adapt standard method [MPP18]

$\rightarrow$ relatively easy result, but lose geometrical interpretation of what happens in strategy space

FTRL in harmonic games admits a constant of motion

Proposition ([Leg+24])

Let $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ be a finite game and consider a vector $m \in \mathbb{R}^N_{++}$ and a fully mixed strategy $q \in \mathcal{X}$. Then the function defined by

$$F_{m,q}(y) := \sum_i m_i \left[h_i(q_i) + h_i^*(y_i) - \langle q_i, y_i \rangle\right]$$

is a constant of motion under FTRL if and only if Γ is harmonic with strategic center (m, q).

FTRL is divergence-free in all games in payoff space

$$\begin{cases} \dot{y}_i &= v_i(x) \\ x_i &= Q_i(y_i) \end{cases} \implies \dot{y}_i = v_i\big(Q(y)\big) \tag{FTRL}$$

$$\frac{d\dot{y}_{i\alpha_i}}{dy_{j\beta_j}} \equiv 0 \quad \text{by multilinearitity of the payoff functions}$$

- By Liouville's theorem, FTRL in payoff space is volume-preserving in all games ;
- the constant of motion can be used to show that FTRL in harmonic games has only bounded orbits;
- Recurrence follows by Poincaré's theorem. □

Bibliography

[Abd+22] Abdou, Joseph et al. "Decomposition of Games: Some Strategic Considerations". In: *Mathematics of Operations Research* 47.1 (Feb. 2022), pp. 176–208.

[Aue+95] Auer, P. et al. "Gambling in a Rigged Casino: The Adversarial Multi-Armed Bandit Problem". In: *Proceedings of IEEE 36th Annual Foundations of Computer Science*. Oct. 1995, pp. 322–331.

[Can+11] Candogan, Ozan et al. "Flows and Decompositions of Games: Harmonic and Potential Games". In: *Mathematics of Operations Research* 36.3 (Aug. 2011), pp. 474–503.

[Fri91] Friedman, Daniel. "Evolutionary Games in Economics". In: *Econometrica* 59.3 (1991), pp. 637–666. JSTOR: 2938222.

[GP20] Gao, Bolin and Pavel, Lacra. "On the Rate of Convergence of Continuous-Time Game Dynamics in N-Player Potential Games". In: *2020 59th IEEE Conference on Decision and Control (CDC)*. Dec. 2020, pp. 1678–1683.

[HCM17] Heliou, Amélie, Cohen, Johanne, and Mertikopoulos, Panayotis. "Learning with Bandit Feedback in Potential Games". In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., 2017.

[HS00] Hofbauer, Josef and Schlag, Karl H. "Sophisticated Imitation in Cyclic Games". In: *Journal of Evolutionary Economics* 10.5 (Sept. 2000), pp. 523–543.

[Leg+24] Legacci, Davide et al. "No-Regret Learning in Harmonic Games: Extrapolation in the Face of Conflicting Interests". In: In proceeding, May 2024.

[LMP24] Legacci, Davide, Mertikopoulos, Panayotis, and Pradelski, Bary. "A Geometric Decomposition of Finite Games: Convergence vs. Recurrence under Exponential Weights". In: *Forty-First International Conference on Machine Learning*. June 2024.

[LW94] Littlestone, N. and Warmuth, M. K. "The Weighted Majority Algorithm". In: *Information and Computation* 108.2 (Feb. 1994), pp. 212–261.

[MPP18] Mertikopoulos, Panayotis, Papadimitriou, Christos, and Piliouras, Georgios. "Cycles in Adversarial Regularized Learning". In: *Proceedings of the 2018 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Proceedings. Society for Industrial and Applied Mathematics, Jan. 2018, pp. 2703–2717.

[MS18] Mertikopoulos, Panayotis and Sandholm, William H. "Riemannian Game Dynamics". In: *Journal of Economic Theory* 177 (2018), pp. 315–364.

[MS96] Monderer, Dov and Shapley, Lloyd S. "Potential Games". In: *Games and Economic Behavior* 14.1 (May 1996), pp. 124–143.

[San10] Sandholm, William H. *Population Games and Evolutionary Dynamics*. MIT press, 2010.

[Sha79] Shahshahani, Siavash. *A New Mathematical Framework for the Study of Linkage and Selection*. American Mathematical Soc., 1979.

[SS06] Shalev-shwartz, Shai and Singer, Yoram. "Convex Repeated Games and Fenchel Duality". In: *Advances in Neural Information Processing Systems*. Vol. 19. MIT Press, 2006.

[TJ78] Taylor, Peter D and Jonker, Leo B. "Evolutionary Stable Strategies and Game Dynamics". In: *Mathematical biosciences* 40.1-2 (1978), pp. 145–156.