GAMES, DYNAMICS & LEARNING

Panayotis Mertikopoulos[1]

joint with

A. Giannou[2] T. Lianeas[2] E. V. Vlatakis-Gkaragkounis[3]

[1]French National Center for Scientific Research (CNRS) & Criteo AI Lab

[2]NTUA

[3]Columbia University

ECE-NTUA – May 14, 2021

GAMES, DYNAMICS & LEARNING

BASIC CONCEPTS

Panayotis Mertikopoulos[1]

joint with

A. Giannou[2]   T. Lianeas[2]   E. V. Vlatakis-Gkaragkounis[3]

[1]French National Center for Scientific Research (CNRS) & Criteo AI Lab

[2]NTUA

[3]Columbia University

ECE-NTUA – May 14, 2021

What's in a game?

A game* is a collection of three basic primitives:

- A set of players $i \in \mathcal{N}$ (computers, neural nets, biological organisms, ...)
- A set of actions A_i per player $i \in \mathcal{N}$
- A payoff function $u_i : A := \prod_j A_j \to \mathbb{R}$ for each player $i \in \mathcal{N}$

Remarks:

- $\mathcal{N}$ can be finite or continuous
- A_i can be finite or continuous

}  Very different models

- Convention for actions:
  - → "A" and "a" for finite (or unspecified)
  - → "$\mathcal{X}$" and "x" for continuous

* "Strategic" or "normal" form

Notation: $\boxed{\begin{array}{l} \Gamma \equiv \Gamma(\mathcal{N}, A, u) \\ \mathcal{G} = \mathcal{G}(\mathcal{N}, \mathcal{X}, u) \end{array}}$ for finite / for continuous

Example 1: Matching Pennies

Two players: Even and Odd

- Each player (secretly) sets a coin to Heads or Tails
- If the coins match, Even wins \$1; otherwise Odd wins \$1

Game-theoretic formulation

- Players: $\mathcal{N} = \{E, O\}$
- Actions: $\mathcal{A}_E = \mathcal{A}_O = \{H, T\}$
- Payoff functions:
  - → $u_E(H,H) = u_E(T,T) = 1$; $u_E(H,T) = u_E(T,H) = -1$
  - → $u_O(H,H) = u_O(T,T) = -1$; $u_O(H,T) = u_O(T,H) = 1$

Payoff Bimatrix:

$$\begin{array}{cc} & \begin{array}{cc} H & T \end{array} \\ \begin{array}{c} H \\ T \end{array} & \begin{pmatrix} (1,-1) & (-1,1) \\ (-1,1) & (-1,-1) \end{pmatrix} \end{array}$$

Terminology: Matrix Games

↳ Here, two-player zero-sum 2×2 game

Example 2: Kelly auctions

N agents seek to share a splittable resource (computing time, bandwidth, produce, ...)

- Each agent has a budget $b_i \geq 0$
- Each agent bids $x_i \in [0, b_i]$
- Once all bids are in, the resource is split proportionally to the bids

$$\boxed{\text{Player } i \leftsquigarrow \frac{x_i}{\sum_j x_j + c}}$$

"entry barrier"

Game-theoretic formulation

- Players: $\mathcal{N} = \{1, \dots, N\}$
- Actions: $\mathcal{X}_i = [0, b_i]$
- Payoff functions: $u_i(x_1, \dots, x_N) = \dfrac{g_i x_i}{c + \sum_{j=1}^N x_j} - x_i$

value per resource unit

cost per unit

Terminology: Continuous game

refers to actions

Example 3: Traffic routing

A city population seeks to commute from home to work

- Population mass: $\mathcal{M}$
- Set of available paths: $\mathcal{P} = \{ \bullet, \bullet, \bullet, \bullet, \dots \}$
- Traffic along a path: $x_p \in [0, M]$, $\sum_{p \in \mathcal{P}} x_p = \mathcal{M}$
- Load on an edge: $w_e = \sum_{p \ni e} x_p$
- Cost of an edge: $c_e(w_e)$ for some increasing function $c_e : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$
- Cost of a path: $c_p(x) = \sum_{e \in p} c_e(w_e)$

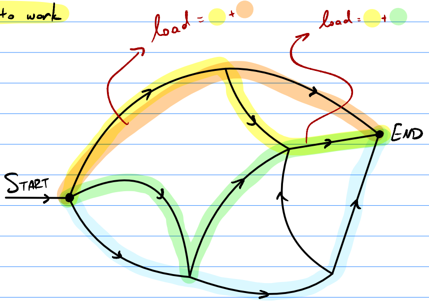

Game-theoretic formulation

- Players: $\mathcal{N} = [0, M]$
- Actions: $A_i = \mathcal{P}$ for all $i \in \mathcal{N}$
- Payoff functions: $u_i(x) = -c_{p_i}(x)$

Is "x" well-defined?

path chosen by i-th player

Cost ( ) = cost($\nearrow$) + cost($\searrow$) + cost($\rightarrow$)

Cost depends only on mass of players choosing a path, not their identity $\leftarrow$ Anonymous Game

A zoo of games

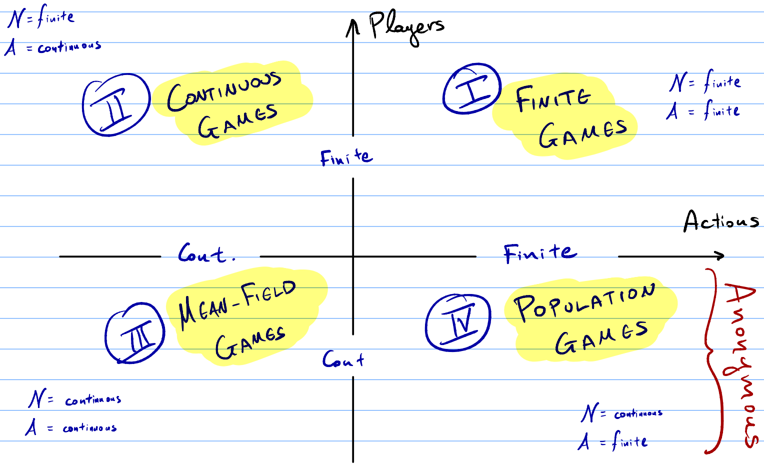

I. Finite Games

Primitives:

- Finite set of players $\mathcal{N} = \{1, \dots, N\}$
- Finite set of actions $\mathcal{A}_i = \{1, \dots, A_i\}$ per player $i \in \mathcal{N}$

also "pure strategies" (more later)

action profile $a = (a_1, \dots, a_N) \equiv (a_i; a_{-i})$

$$\mathcal{A} := \prod_j \mathcal{A}_j = \mathcal{A}_i \times \mathcal{A}_{-i} := \prod_{j \neq i} \mathcal{A}_j$$

[all players] = [i-th player] × [other players]

- Payoff functions $u_i : \mathcal{A} \to \mathbb{R}$, $i = 1, \dots, N$

Assumptions:

- Sometimes $u_i \in [-1, 1]$ or $[0, 1]$ otherwise NONE

I. Finite Games & their mixed extensions

Mixed strategies:

Players can mix their actions by playing a probability distribution x_i on $\mathcal{A}_i$

- $x_{ia_i} \longrightarrow$ Player $i \in \mathcal{N}$ selects action $a_i \in \mathcal{A}_i$ w/ prob. x_{ia_i}
- $x_{ia_i} \geq 0$, $\sum_{a_i \in \mathcal{A}_i} x_{ia_i} = 1$
- **Probability simplex:**

$$\boxed{\mathcal{X}_i := \Delta(\mathcal{A}_i) \equiv \{x_i \in \mathbb{R}^{\mathcal{A}_i} : x_{ia_i} \geq 0, \ \sum_{a_i \in \mathcal{A}_i} x_{ia_i} = 1\}}$$

- **Mixed strategy profile:** $x = (x_1, \dots, x_N) \equiv (x_i ; x_{-i})$

$$\mathcal{X} := \prod_j \mathcal{X}_j \quad = \quad \mathcal{X}_i \quad \times \quad \mathcal{X}_{-i} := \prod_{j \neq i} \mathcal{X}_j$$

[all players] = [i-th player] × [other players]

- **Mixed / Expected payoffs:**

$$\boxed{\begin{aligned} u_i(x) &= \sum_{a_1 \in \mathcal{A}_1} \cdots \sum_{a_N \in \mathcal{A}_N} x_{1a_1} \cdots x_{Na_N} \, u_i(a_1, \dots a_N) \\ u_i(x_i ; x_{-i}) &= \sum_{a_i \in \mathcal{A}_i} \sum_{a_{-i} \in \mathcal{A}_{-i}} x_{ia_i} \cdot x_{-i,a_{-i}} \, u_i(a_i ; a_{-i}) \end{aligned}}$$

$x_{-i,a_{-i}} := \prod_{j \neq i} x_{ja_j}$: mixed profile of "other" players

- **Notation:** $\mathcal{G} = \Delta(\Gamma)$ or $\overline{\Gamma}$

I. Finite Games: Examples

Example: Matching Pennies:

- Players: $\mathcal{N} = \{E, O\}$
- Actions: $\mathcal{A}_E = \mathcal{A}_O = \{H, T\}$
- Payoff functions:
  - $\to u_E(H,H) = u_E(T,T) = 1$; $u_E(H,T) = u_E(T,H) = -1$
  - $\to u_O(H,H) = u_O(T,T) = -1$; $u_O(H,T) = u_O(T,H) = 1$

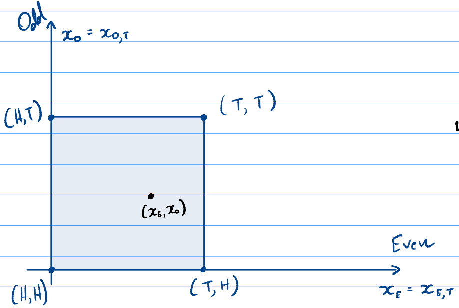

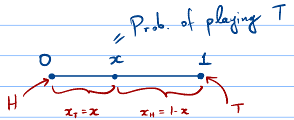

$$
\begin{aligned}
u_E(x_E, x_O) = {} & x_E\, x_O \; \cancelto{1}{u_E(T,T)} \\
+ {} & x_E (1-x_O) \; \cancelto{-1}{u_E(T,H)} \\
+ {} & (1-x_E)\, x_O \; \cancelto{-1}{u_E(H,T)} \\
+ {} & (1-x_E)(1-x_O) \; \cancelto{1}{u_E(H,H)}
\end{aligned}
$$

I. Finite Games: Examples

Exercise: Rock-Paper-Scissors

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_1 = \mathcal{A}_2 = \{R, P, S\}$
- Payoff functions:

  → $u_1(R,R) = u_1(P,P) = u_1(S,S) = 0$; $u_1(R,P) = u_1(P,S) = u_1(S,R) = -1$;

  $u_1(P,R) = u_1(R,S) = u_1(S,P) = 1$

  → $u_2 = -u_1$

- Write out mixed extension of RPS (mixed strategy representation + mixed payoffs)
- Propose strategy for Player 2 if Player 1 plays (R, P, S) w/ prob a) $(1/3, 1/3, 1/3)$

  b) $(1/3, 2/3, 0)$

  c) $(1/2, 1/2, 0)$

II. Continuous Games: Basics

Primitives:

- Finite set of players $\mathcal{N} = \{1, \dots, N\}$
- Continuous set of actions $\mathcal{X}_i$ per player $i \in \mathcal{N}$

only pure strategies here

action profile $= (x_1, \dots, x_N) \equiv (x_i ; x_{-i})$

$\mathcal{X} := \prod_j \mathcal{X}_j = \mathcal{X}_i \times \mathcal{X}_{-i} := \prod_{j \neq i} \mathcal{X}_j$

[all players] = [i-th player] × [other players]

- Payoff functions $u_i : \mathcal{X} \to \mathbb{R}$, $i = 1, \dots, N$

Assumptions:

- $\mathcal{X}_i$ = closed, convex subset of ambient vector space $\mathcal{V}_i \cong \mathbb{R}^{d_i}$ w/ norm $\|\cdot\|$
- $u_i(x_i ; x_{-i})$ concave in x_i for all $x_{-i} \in \mathcal{X}_{-i}$, $i \in \mathcal{N}$ ⟶ concave game
- $u_i(x)$ C^1-smooth in x for all $x \in \mathcal{X}$, $i \in \mathcal{N}$ ⟶ smooth game

ATTN: do not confuse w/ (λ, μ)-smooth games

II. Continuous Games: Examples

Example: Mixed Extensions Revisited

- Finite set of players $\mathcal{N} = \{1, \dots, N\}$
- Action sets $\mathcal{X}_i = \Delta(\mathcal{A}_i) \subseteq \mathbb{R}^{\mathcal{A}_i}$ for some finite set $\mathcal{A}_i$
- Payoff functions:

$$\boxed{u_i(x) = \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} x_{1\alpha_1} \cdots x_{N\alpha_N} \, u_i(\alpha_1, \dots \alpha_N)}$$

- Verify assumptions:
  - → $\mathcal{X}_i$ closed, compact (why?)
  - → $u_i(x_i; x_{-i})$ $\underbrace{\text{linear}}_{\text{concave}}$ in x_i, $\underbrace{\text{multilinear}}_{\text{smooth}}$ in $x = (x_1, \dots, x_N)$

- Mixed extensions of finite games are smooth concave games

Exercise: Verify that Kelly auctions are smooth, concave games.

III. Population Games: Basics

Primitives:

- Continuous population of players $i \in \mathcal{I} = [0,1]$, endowed w/ Lebesgue measure μ $[\mu([a,b]) = b-a]$
- Shared finite set of actions $\mathcal{A} = \{1, \dots, A\}$
- Strategy profile = measurable assignment $\chi : \mathcal{I} \to \mathcal{A}$

  $\chi(i) \in \mathcal{A}$ = action choice of player $i \in \mathcal{I}$
- Set of "a-strategists": $\chi^{-1}(a) = \{i \in \mathcal{I} : \chi(i) = a\}$

- Population state: $x = \chi \# \mu = \mu \circ \chi^{-1}$

  $\boxed{x_a = \mu\{i \in \mathcal{I} : \chi(i) = a\} = \text{mass of players playing } a \in \mathcal{A}}$

  Set of population states $\boxed{\mathcal{X} = \{x \in \mathbb{R}^{\mathcal{A}} : x_a \geq 0, \sum_{a \in \mathcal{A}} x_a = 1\}}$

- __Anonymity:__ player payoffs only depend ("factor through") the state of the population $x \in \mathcal{X}$

  Payoff functions $u_a : \mathcal{X} \to \mathbb{R}$: $u_a(x)$ = payoff to a-strategists in pop. state $x \in \mathcal{X}$

III. Population Games: Examples

Example 1: Nonatomic Congestion Games

- Player population: $\mathcal{I} = [0, M]$ (renormalize)
- Actions: network routes $p \in \mathcal{P}$ joining origin to destination
- Payoff functions: $u_p(x) = -c_p(x)$ defined as before

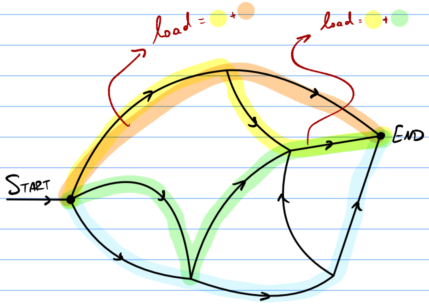

$\text{cost}(\;) = \text{cost}(\nearrow) + \text{cost}(\searrow) + \text{cost}(\rightarrow)$

III. Population Games: Examples

Example 2: Single-population random matching

- **Player population:** $\mathcal{J} = [0,1]$
- **Given:** symmetric 2-player finite game Γ
  - Common action set $\mathcal{A} = \{1, \dots, A\}$
  - Symmetric payoff functions: $u_1(\alpha_1, \alpha_2) = u(\alpha_1, \alpha_2) = u_2(\alpha_2, \alpha_1)$
  - payoffs given by matrix, not bimatrix

- **Random matching:** two players selected uniformly at random from the population and face each other in Γ

- **Mean payoff to α-strategists:** $$\boxed{u_\alpha(x) = \sum_{\beta \in \mathcal{A}} x_\beta \, u(\alpha, \beta)}$$

- **Mean population payoff:** $u(x) = \sum_{\alpha \in \mathcal{A}} x_\alpha u_\alpha(x) = \sum_{\alpha,\beta \in \mathcal{A}} x_\alpha \, u(\alpha,\beta) \, x_\beta$
  quadratic in x

II. Population Games: Several Populations

Primitives:

think of "player types"

- Several populations of players $i = 1, \dots, N$
- Shared finite set of actions $\mathcal{A}_i = \{1, \dots, A_i\}$ per population

- Population states: x_{ia_i} = mass of i-type players playing $a_i \in \mathcal{A}_i$

  $x_i = (x_{ia_i})_{a_i \in \mathcal{A}_i}$ = mass distribution = state of i-th population

  $x = (x_1, \dots, x_N)$ = collective population state

- Anonymity: player payoffs only depend ("factor through") the state of the population $x \in \mathcal{X}$

  Payoff functions $u_{ia_i}: \mathcal{X} \to \mathbb{R}$: $u_{ia_i}(x)$ = payoff to a_i-strategists in pop. state $x \in \mathcal{X}$

Assumptions: u_{ia_i} is Lipschitz continuous for all $a_i \in \mathcal{A}_i$, $i \in \mathcal{N}$

III. Population Games: Examples

Example 3: Multi-population random matching

- **Given:** N player populations (unit mass)
- **Given:** finite N-player game $\Gamma = (\mathcal{N}, \mathcal{A}, u)$
  - different action sets $\mathcal{A}_i = \{1, \dots, A_i\}$ per player / player type
  - payoffs given by polymatrix

- **Random matching:** N players selected uniformly at random, one from each population, matched against each other in Γ

- Mean payoff to α_i-strategists of i-th population:

$$\boxed{u_{i\alpha_i}(x) = \sum_{\beta_1 \in \mathcal{A}_1} \cdots \sum_{\beta_N \in \mathcal{A}_N} x_{1,\beta_1} \cdots \delta_{\alpha_i \beta_i} \cdots x_{N,\beta_N} \, u_i(\beta_1, \dots, \beta_N)}$$

- Mean population payoff: $u_i(x) = \sum_{\alpha_i \in \mathcal{A}_i} x_{i\alpha_i} u_{i\alpha_i}(x) = \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} x_{1\alpha_1} \cdots x_{N,\alpha_N} \, u_i(\alpha_1, \dots, \alpha_N)$

multilinear in x

Relations between games

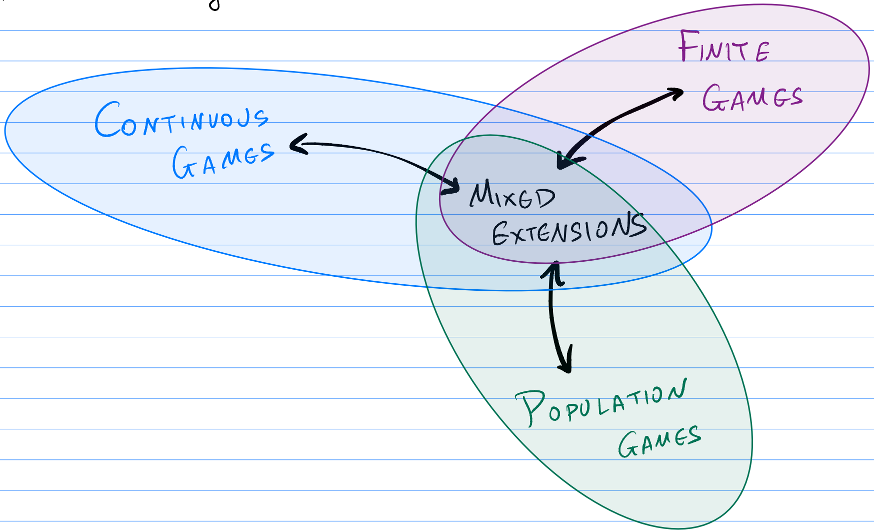

A zoo of games