GAMES, DYNAMICS & LEARNING

2. POPULATION GAMES AND EVOLUTIONARY DYNAMICS

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 21, 2021

Outline

Bird's eye view

Last time:

- Games, examples, taxonomy,...
- Congestion games (atomic / nonatomic, splittable / non-splittable,...)

Moving forward: playing day-by-day

- **Lecture 2:** population games $\leftrightsquigarrow$ evolutionary dynamics
- **Lecture 3:** finite games $\leftrightsquigarrow$ multi-armed bandits
- **Lecture 4:** continuous games $\leftrightsquigarrow$ online convex optimization

Bird's eye view

Last time:

- Games, examples, taxonomy,…
- Congestion games (atomic / nonatomic, splittable / non-splittable,…)

Moving forward: playing day-by-day

- **Lecture 2: population games ⇝ evolutionary dynamics**
- **Lecture 3:** finite games ⇝ multi-armed bandits
- **Lecture 4:** continuous games ⇝ online convex optimization

Caveats

- **Big picture:** Focus on concepts + selected deep dives
- **Notation:** losses ("ℓ") ⇝ utilities ("u"),…; pure strategies ⇝ actions; etc.

Today: Playing day after day

A typical **online decision process:**

repeat

At each epoch $t \geq 0$

Choose **action** [focal player]

Incur **loss** / Receive **reward** [depends on context]

Get **feedback** [depends on context]

until end

Key considerations

- **Time:** continuous or discrete?
- **Players:** continuous or discrete?
- **Actions:** continuous or discrete?
- **Payoffs / Losses:** determined by other players or "Nature"?
- **Feedback:** full info / observation? payoff-based?

Outline

Population games

- **Players:** continuous, nonatomic *populations* $i = 1, \dots, N$ [species, types,...]
- **Actions:** finite *action set* $\mathcal{A}_i$ per population [phenotypes, routes,...]
- **Payoffs:** depend only on the players' *distribution* [anonymity]
  - *Population shares*

$$x_{ia_i} = \text{relative frequency of } a_i \in \mathcal{A}_i \text{ in population } i$$

  - *Population states*

$$x_i = (x_{ia_i})_{a_i \in \mathcal{A}_i} \in \mathcal{X}_i \coloneqq \Delta(\mathcal{A}_i)$$

$$x = (x_1, \dots, x_N) \in \mathcal{X} \coloneqq \prod_i \mathcal{X}_i$$

  - *Payoff functions* $u_{ia_i} \colon \mathcal{X} \to \mathbb{R}$

$$u_{ia_i}(x) = \text{payoff to } a_i \in \mathcal{A}_i \text{ when the population is at state } x \in \mathcal{X}$$

  - *Mean population payoff*

$$u_i(x) = \sum_{a_i \in \mathcal{A}_i} x_{ia_i} u_{ia_i}(x)$$

Examples and more

Example 1: Multi-population random matching [Maynard Smith and Price, 1973]

- Given: finite N-player game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ [base game]
- Given: N player populations, each with action set $\mathcal{A}_i$
- During play:
  - Players drawn uniformly at random from each population
  - Drawn players matched to play Γ [random matching]
  - Mean payoffs

$$u_{ia_i}(x) = \sum_{a'_1 \in \mathcal{A}_1} \cdots \sum_{a'_N \in \mathcal{A}_N} x_{1,a'_1} \cdots \delta_{a'_i a_i} \cdots x_{N,a'_N} u_i(a'_1, \dots, a'_N)$$

  - **Caveat** Single-population matching is different (quadratic) [why?]

Examples and more

Example 1: Multi-population random matching [Maynard Smith and Price, 1973]

- Given: finite N-player game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ [base game]
- Given: N player populations, each with action set $\mathcal{A}_i$
- During play:
  - Players drawn uniformly at random from each population
  - Drawn players matched to play Γ [random matching]
  - Mean payoffs

$$u_{ia_i}(x) = \sum_{a'_1 \in \mathcal{A}_1} \cdots \sum_{a'_N \in \mathcal{A}_N} x_{1,a'_1} \cdots \delta_{a'_i a_i} \cdots x_{N,a'_N} u_i(a'_1, \dots, a'_N)$$

  - **Caveat** Single-population matching is different (quadratic) [why?]

Example 2: Nonatomic congestion games ("playing the field")

- See Lecture 1

For more: Weibull [1995], Hofbauer and Sigmund [1998], Sandholm [2010, 2015]

Relations between classes

Porous boundaries:

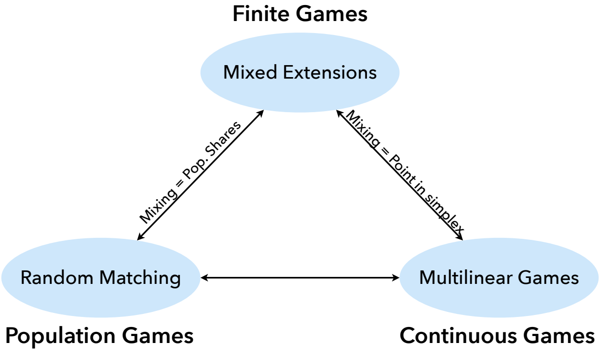

Important

- Random matching $\subsetneq$ population games of interest
- (Multi)Linear games $\subsetneq$ continuous games of interest

Solution concepts

- **Dominated strategies:** $a_i \in \mathcal{A}_i$ is *dominated* by $a_i' \in \mathcal{A}_i$ ($a_i \prec a_i'$) if

$$u_{ia_i}(x) < u_{ia_i'}(x) \qquad \text{for all } x \in \mathcal{X}$$

[Mixed version: $p_i \prec q_i \iff u_i(p_i; x_{-i}) < u_i(q_i; x_{-i})$ for all $x \in \mathcal{X}$]

The Prisoner's Dilemma

Two villains, Robin and Charlie, are caught

→ "Rat on your compadre, and you go free"

→ "Stay loyal, face 10 years in jail"

	Charlie silent	Charlie betrays
Robin silent	(-1, -1)	(-10, 0)
Robin betrays	(0, -10)	(-2, -2)

What is the outcome?

silent < betray

Iteratively Dominated Strategies

<u>Another example:</u>

(9,4)	(5,3)	(3,2)
(0,1)	(4,6)	(6,0)
(2,1)	(3,5)	(2,8)

Step 1 (eliminate row 3), Step 2 (eliminate column 3), Step 3 (eliminate row 2), Step 4 (eliminate (5,3)); (9,4) boxed.

What is the outcome of this game?

→ Iteratively dominated strategy
  ↳ Dominated strategy in the subgame resulting from the elimination of dominated strategies

→ Dominance solvable game:
  ↳ Elimination of iteratively dominated strategies leads to a singleton

Solution concepts

- **Dominated strategies:** $a_i \in \mathcal{A}_i$ is *dominated* by $a_i' \in \mathcal{A}_i$ ($a_i \prec a_i'$) if

$$u_{ia_i}(x) < u_{ia_i'}(x) \qquad \text{for all } x \in \mathcal{X}$$

[Mixed version: $p_i \prec q_i \iff u_i(p_i; x_{-i}) < u_i(q_i; x_{-i})$ for all $x \in \mathcal{X}$]

- **Nash equilibrium:** no player has an incentive to switch strategies

$$u_{ia_i}(x^*) \geq u_{ia_i'}(x^*) \qquad \text{for all } a_i, a_i' \in \mathcal{A}_i \text{ with } x^*_{ia_i} > 0$$

  - *Support* of x^*: $\operatorname{supp}(x^*) = \{(a_1, \dots, a_N) \in \mathcal{A} : x^*_{ia_i} > 0 \text{ for all } i\}$
  - *Interior / Full support equilibria:* $\operatorname{supp}(x^*) = \mathcal{A}$
  - *Pure equilibria:* $\operatorname{supp}(x^*) =$ singleton
  - *Strict equilibria:* ">" instead of "≥" when $a_i' \notin \operatorname{supp}(x_i^*)$

Solution concepts

- **Dominated strategies:** $a_i \in \mathcal{A}_i$ is *dominated* by $a_i' \in \mathcal{A}_i$ ($a_i \prec a_i'$) if

$$u_{ia_i}(x) < u_{ia_i'}(x) \qquad \text{for all } x \in \mathcal{X}$$

[Mixed version: $p_i \prec q_i \iff u_i(p_i; x_{-i}) < u_i(q_i; x_{-i})$ for all $x \in \mathcal{X}$]

- **Nash equilibrium:** no player has an incentive to switch strategies

$$u_{ia_i}(x^*) \geq u_{ia_i'}(x^*) \qquad \text{for all } a_i, a_i' \in \mathcal{A}_i \text{ with } x^*_{ia_i} > 0$$

  - *Support* of x^*: $\operatorname{supp}(x^*) = \{(a_1, \dots, a_N) \in \mathcal{A} : x^*_{ia_i} > 0 \text{ for all } i\}$
  - *Interior / Full support equilibria:* $\operatorname{supp}(x^*) = \mathcal{A}$
  - *Pure equilibria:* $\operatorname{supp}(x^*)$ = singleton
  - *Strict equilibria:* ">" instead of "≥" when $a_i' \notin \operatorname{supp}(x_i^*)$

- **Examples:**
  - *Rock-Paper-Scissors*: unique, full support equilibrium
  - *Prisoner's dilemma*: dominance-solvable, one strict equilibrium
  - *Battle of the Sexes*: one full support equilibrium; two strict equilibria

Go-to example: Rock-Paper-Scissors

- Players: $\mathcal{N} = \{1, 2\}$.

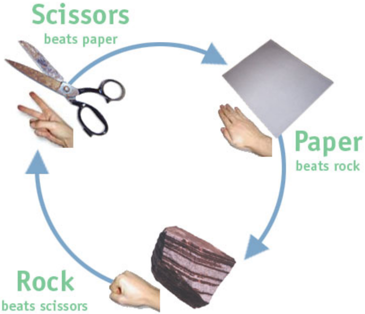

Go-to example: Rock-Paper-Scissors

- Players: $\mathcal{N} = \{1, 2\}$.
- Actions: $\mathcal{A}_i = \{\mathsf{R}, \mathsf{P}, \mathsf{S}\}, i = 1, 2$.
- Payoff matrix (win 1, lose −1, tie 0):

$A =$

	R	P	S
R	0	−1	1
P	1	0	−1
S	−1	1	0

- Mixed strategies: $x_i \in \Delta\{\mathsf{R}, \mathsf{P}, \mathsf{S}\}$.
- Payoff functions (multi-population):

$$u_1(x) = -u_2(x) = x_1^\top A x_2$$

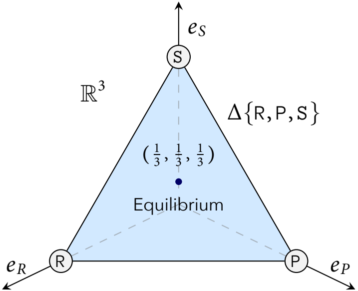

Multi-population matching ⇝ Mixed extension ⇝ Multilinear game

Go-to example: Rock-Paper-Scissors

- Players: $\mathcal{N} = \{1, 2\}$.
- Actions: $\mathcal{A}_i = \{\mathrm{R}, \mathrm{P}, \mathrm{S}\}, i = 1, 2$.
- Payoff matrix (win 1, lose −1, tie 0):

$A =$

	R	P	S
R	0	−1	1
P	1	0	−1
S	−1	1	0

- Mixed strategies: $x_i \in \Delta\{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$.
- Payoff functions (single-population):

$$u(x) = x^\top A x$$

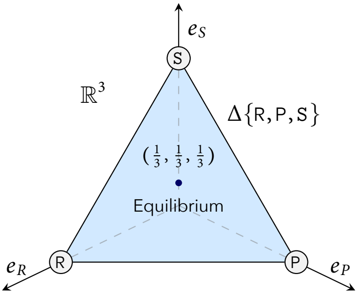

Single-population matching ⇜̸⇝ Mixed extensions ⇜⇝ Multilinear games

Rest of this lecture

Are game-theoretic solution concepts consistent with evolutionary models?

- Evolutionary models $\rightsquigarrow$ dynamical systems (Lotka-Volterra, replicator, etc.)
- Do dominated strategies become extinct?
- Is equilibrium play stable/attracting?
- …

Outline

A biologist's viewpoint

- Populations are *species*, strategies are *phenotypes:*

$$z_{ia_i} = \text{absolute population mass of type } a_i \in \mathcal{A}_i$$
$$z_i = \sum_{a_i} z_{ia_i} = \text{absolute population mass of } i\text{-th species}$$

A biologist's viewpoint

- Populations are *species,* strategies are *phenotypes:*

$$z_{ia_i} = \text{absolute population mass of type } a_i \in \mathcal{A}_i$$
$$z_i = \sum\nolimits_{a_i} z_{ia_i} = \text{absolute population mass of } i\text{-th species}$$

- Utilities measure *fecundity / reproductive fitness:*

$$u_{ia_i}(x) = \text{per capita growth rate of type } a_i$$

- Population evolution:

$$\dot{z}_{ia_i} = z_{ia_i} u_{ia_i}$$

A biologist's viewpoint

- Populations are *species,* strategies are *phenotypes:*

$$z_{ia_i} = \text{absolute population mass of type } a_i \in \mathcal{A}_i$$
$$z_i = \textstyle\sum_{a_i} z_{ia_i} = \text{absolute population mass of } i\text{-th species}$$

- Utilities measure *fecundity / reproductive fitness:*

$$u_{ia_i}(x) = \text{per capita growth rate of type } a_i$$

- Population evolution:

$$\dot{z}_{ia_i} = z_{ia_i} u_{ia_i}$$

- Evolution of population shares ($x_{ia_i} = z_{ia_i}/z_i$):

$$\dot{x}_{ia_i} = \frac{d}{dt}\frac{z_{ia_i}}{z_i} = \frac{\dot{z}_{ia_i} z_i - z_{ia_i}\sum_{a_i'} \dot{z}_{ia_i'}}{z_i^2} = \frac{z_{ia_i}}{z_i} u_{ia_i} - \frac{z_{ia_i}}{z_i}\sum_{a_i'} \frac{z_{ia_i'}}{z_i} u_{ia_i'}$$

Replicator dynamics [Taylor and Jonker, 1978]

$$\dot{x}_{ia_i} = x_{ia_i}\left[u_{ia_i}(x) - u_i(x)\right] \tag{RD}$$

An economist's viewpoint

- Agents receive **revision opportunities** to switch strategies

$$\rho_{aa'}(x) = \text{conditional switch rate from } a \text{ to } a'$$

[**NB:** dropping player index for simplicity]

An economist's viewpoint

- Agents receive **revision opportunities** to switch strategies

$$\rho_{aa'}(x) = \text{conditional switch rate from } a \text{ to } a'$$

[NB: dropping player index for simplicity]

- **Pairwise proportional imitation:**

$$\rho_{aa'}(x) = x_{a'}[u_{a'}(x) - u_a(x)]_+$$

[Imitate with probability proportional to excess payoff (Helbing, 1992; Schlag, 1998)]

- Inflow/outflow:

$$\textit{Incoming toward } a = \sum_{a'} \operatorname{mass}(a' \rightsquigarrow a) = \sum_{a' \in \mathcal{A}} x_{a'} \rho_{a'a}(x)$$

$$\textit{Outgoing from } a = \sum_{a'} \operatorname{mass}(a \rightsquigarrow a') = x_a \sum_{a' \in \mathcal{A}} \rho_{aa'}(x)$$

Conditional switch: $\rho_{\alpha\beta} = x_\beta [u_\beta - u_\alpha]_+$

Inflow to α = $\sum_{\beta\neq\alpha} x_\beta \rho_{\beta\alpha}(x) = \sum_\beta x_\beta x_\alpha [u_\alpha - u_\beta]_+$

Outflow from α = $\sum_{\beta\neq\alpha} x_\alpha \rho_{\alpha\beta}(x) = \sum_\beta x_\alpha x_\beta [u_\beta - u_\alpha]_+$

Net flow $\dot{x}_\alpha = x_\alpha \left[\sum_\beta x_\beta [u_\alpha - u_\beta]_+ - \sum_\beta x_\beta [u_\beta - u_\alpha]_+ \right]$

"+": $u_\alpha \geq u_\beta$

"−": $u_\alpha \leq u_\beta$

$$= x_\alpha \left[\sum_\beta^{+} x_\beta [u_\alpha - u_\beta] + \sum_\beta^{-} x_\beta [u_\alpha - u_\beta] \right]$$

$$= x_\alpha \sum_\beta x_\beta [u_\alpha - u_\beta]$$

$$= x_\alpha \left[u_\alpha - \sum_\beta x_\beta u_\beta \right] = x_\alpha [u_\alpha(x) - u(x)]$$

An economist's viewpoint

- Agents receive **revision opportunities** to switch strategies

$$\rho_{aa'}(x) = \text{conditional switch rate from } a \text{ to } a'$$

[**NB:** dropping player index for simplicity]

- **Pairwise proportional imitation:**

$$\rho_{aa'}(x) = x_{a'}[u_{a'}(x) - u_a(x)]_+$$

[Imitate with probability proportional to excess payoff (Helbing, 1992; Schlag, 1998)]

- Inflow/outflow:

$$\textit{Incoming toward } a = \sum_{a'} \text{mass}(a' \rightsquigarrow a) = \sum_{a' \in \mathcal{A}} x_{a'} \rho_{a'a}(x)$$

$$\textit{Outgoing from } a = \sum_{a'} \text{mass}(a \rightsquigarrow a') = x_a \sum_{a' \in \mathcal{A}} \rho_{aa'}(x)$$

- Detailed balance:

$$\dot{x}_a = \text{inflow}_a(x) - \text{outflow}_a(x) = \cdots = x_a[u_a(x) - u(x)] \qquad \text{(RD)}$$

A learning viewpoint

Evolution of mixed strategies in a finite game:

- Agents record cumulative payoff of each strategy

$$y_a(t) = \int_0^t u_a(\tau)\, d\tau$$

$\Longrightarrow$ *propensity* of choosing a strategy [Littlestone and Warmuth, 1994; Vovk, 1995]

- Choice probabilities $\rightsquigarrow$ exponentially proportional to propensity scores

$$x_a(t) = \frac{\exp(y_a(t))}{\sum_{a'} \exp(y_{a'}(t))}$$

$$x_\alpha = \frac{\exp(y_\alpha)}{\sum_\beta \exp(y_\beta)}$$

$$\dot{x}_\alpha = \frac{\frac{d}{dt}\exp(y_\alpha) \overbrace{\sum_\beta \exp(y_\beta)}^{Z} - \exp(y_\alpha)\frac{d}{dt}\overbrace{\sum_\beta \exp(y_\beta)}^{Z}}{\big[\underbrace{\sum_\beta \exp(y_\beta)}_{Z}\big]^2}$$

$$= \frac{\overbrace{\dot{y}_\alpha}^{= u_\alpha} \exp(y_\alpha) Z - \exp(y_\alpha) \sum_\beta \overbrace{\dot{y}_\beta}^{= u_\beta} \exp(y_\beta)}{Z^2}$$

$$= x_\alpha u_\alpha - x_\alpha \sum_\beta x_\beta u_\beta$$

$$= x_\alpha \left[u_\alpha(x) - u(x)\right]$$

A learning viewpoint

Evolution of mixed strategies in a finite game:

- Agents record cumulative payoff of each strategy

$$y_a(t) = \int_0^t u_a(\tau)\, d\tau$$

$\Longrightarrow$ *propensity* of choosing a strategy [Littlestone and Warmuth, 1994; Vovk, 1995]

- Choice probabilities $\rightsquigarrow$ exponentially proportional to propensity scores

$$x_a(t) = \frac{\exp(y_a(t))}{\sum_{a'} \exp(y_{a'}(t))}$$

- Evolution of mixed strategies [Rustichini, 1999; Hofbauer et al., 2009]

$$\dot{x}_a = \cdots = x_a[u_a(x) - u(x)] \tag{RD}$$

[**Check:** verify computation]

Basic properties

Different viewpoints, same dynamics:

$$\dot{x}_{ia_i} = x_{a_i}[u_{ia_i}(x) - u_i(x)] \tag{RD}$$

[**NB:** all viewpoints will be useful later]

Structural properties

[Weibull, 1995; Hofbauer and Sigmund, 1998]

- **Well-posed:** every initial condition $x \in \mathcal{X}$ admits unique solution trajectory $x(t)$ that exists for all time

  [Assuming u_i is Lipschitz]

- **Consistent:** $x(t) \in \mathcal{X}$ for all $t \geq 0$

  [Assuming $x(0) \in \mathcal{X}$]

- **Faces are forward invariant** ("strategies breed true"):

$$x_{ia_i}(0) > 0 \iff x_{ia_i}(t) > 0 \quad \text{for all } t \geq 0$$
$$x_{ia_i}(0) = 0 \iff x_{ia_i}(t) = 0 \quad \text{for all } t \geq 0$$

Outline

Dynamics and rationality

Are game-theoretic solution concepts consistent with evolutionary models?

Dynamics and rationality

Are game-theoretic solution concepts consistent with the replicator dynamics?

- Do dominated strategies die out in the long run?
- Are Nash equilibria stationary?
- Are they stable? Are they attracting?
- Do the replicator dynamics always converge?
- What other behaviors can we observe?
- ...

Phase portraits

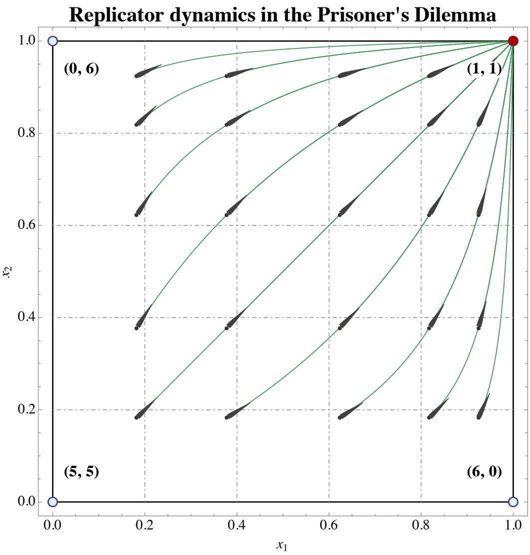

Phase portraits

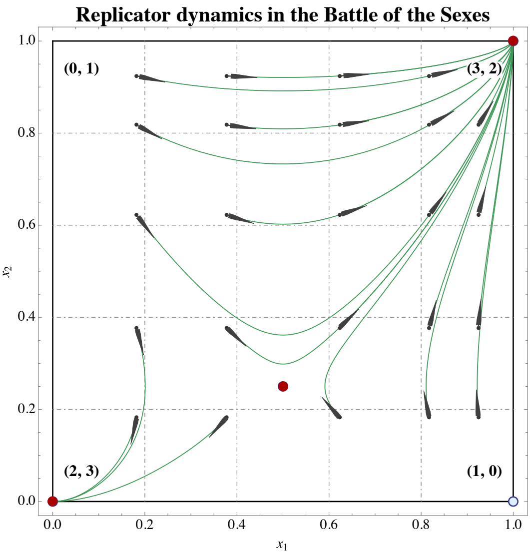

Phase portraits

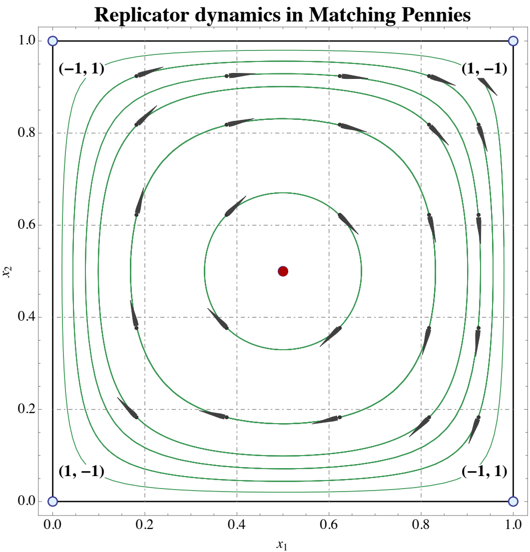

Phase portraits

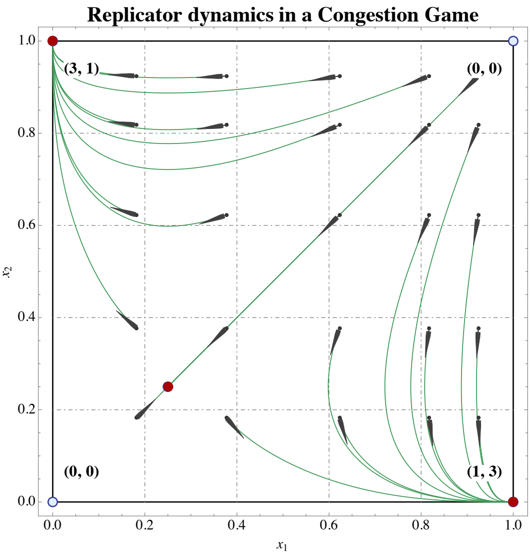

Dominated strategies

Suppose $a_i \in \mathcal{A}_i$ is *dominated* by $a_i' \in \mathcal{A}_i$

- Consistent payoff gap:

$$u_{ia_i}(x) \leq u_{ia_i'}(x) - \varepsilon \quad \text{for some } \varepsilon > 0$$

Dominated strategies

Suppose $a_i \in \mathcal{A}_i$ is *dominated* by $a_i' \in \mathcal{A}_i$

- Consistent payoff gap:

$$u_{ia_i}(x) \leq u_{ia_i'}(x) - \varepsilon \quad \text{for some } \varepsilon > 0$$

- Consistent difference in scores:

$$y_{ia_i}(t) = \int_0^t u_{ia_i}(x)\,d\tau \leq \int_0^t [u_{ia_i'}(x) - \varepsilon]\,d\tau = y_{ia_i'}(t) - \varepsilon t$$

- Consistent difference in choice probabilities

$$\frac{x_{ia_i}(t)}{x_{ia_i'}(t)} = \frac{\exp(y_{ia_i}(t))}{\exp(y_{ia_i'}(t))} \leq \exp(-\varepsilon t)$$

Dominated strategies

Suppose $a_i \in \mathcal{A}_i$ is *dominated* by $a_i' \in \mathcal{A}_i$

- Consistent payoff gap:

$$u_{ia_i}(x) \leq u_{ia_i'}(x) - \varepsilon \quad \text{for some } \varepsilon > 0$$

- Consistent difference in scores:

$$y_{ia_i}(t) = \int_0^t u_{ia_i}(x)\, d\tau \leq \int_0^t [u_{ia_i'}(x) - \varepsilon]\, d\tau = y_{ia_i'}(t) - \varepsilon t$$

- Consistent difference in choice probabilities

$$\frac{x_{ia_i}(t)}{x_{ia_i'}(t)} = \frac{\exp(y_{ia_i}(t))}{\exp(y_{ia_i'}(t))} \leq \exp(-\varepsilon t)$$

- **Dominated strategies become extinct** [Samuelson and Zhang, 1992]

$$\lim_{t\to\infty} x_{ia_i}(t) = 0 \quad \text{whenever } a_i \text{ is dominated}$$

[**Check #1:** extend to *iteratively / mixed* dominated strategies]

[**Check #2:** what about *weakly* dominated strategies?]

Stationarity of equilibria

Nash equilibrium: $u_{ia_i}(x^*) \geq u_{ia_i'}(x^*)$ for all $a_i, a_i' \in \mathcal{A}_i$ with $x^*_{ia_i} > 0$

- Supported strategies have equal payoffs:

$$u_{ia_i}(x^*) = u_{ia_i'}(x^*) \quad \text{for all } a_i, a_i' \in \operatorname{supp}(x_i^*)$$

- Mean payoff equal to equilibrium payoff:

$$u_i(x^*) = u_{ia_i}(x^*) \quad \text{for all } a_i \in \operatorname{supp}(x_i^*)$$

Stationarity of equilibria

Nash equilibrium: $u_{ia_i}(x^*) \geq u_{ia_i'}(x^*)$ for all $a_i, a_i' \in \mathcal{A}_i$ with $x^*_{ia_i} > 0$

- Supported strategies have equal payoffs:

$$u_{ia_i}(x^*) = u_{ia_i'}(x^*) \quad \text{for all } a_i, a_i' \in \operatorname{supp}(x_i^*)$$

- Mean payoff equal to equilibrium payoff:

$$u_i(x^*) = u_{ia_i}(x^*) \quad \text{for all } a_i \in \operatorname{supp}(x_i^*)$$

- Replicator field vanishes at equilibria:

$$x^*_{ia_i}[u_{ia_i}(x^*) - u_i(x^*)] = 0 \quad \text{for all } a_i \in \mathcal{A}_i$$

- **Nash equilibria are stationary:**

$$x(0) = x^* \iff x(t) = x^* \text{ for all } t \geq 0$$

- **The converse does not hold** (never used inequality)

[**Check:** All vertices are stationary – general statement?]

Stability

Are all stationary points created equal?

Definition

x^* is **(Lyapunov) stable** if, for every neighborhood U of x^* in $\mathcal{X}$, there exists a neighborhood U' of x^* such that

$$x(0) \in U' \implies x(t) \in U \quad \text{for all } t \geq 0$$

[Trajectories that start close to x^* remain close for all time]

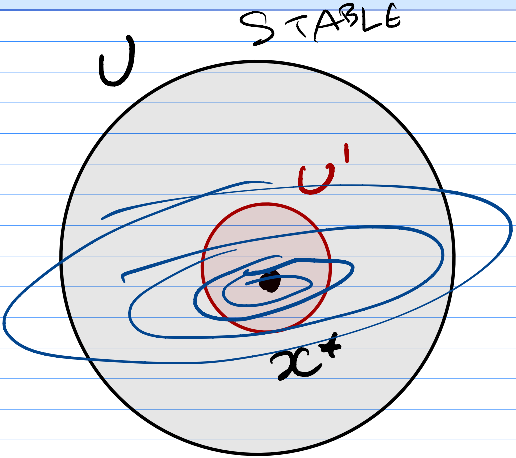
STABLE
U
U'
x^*

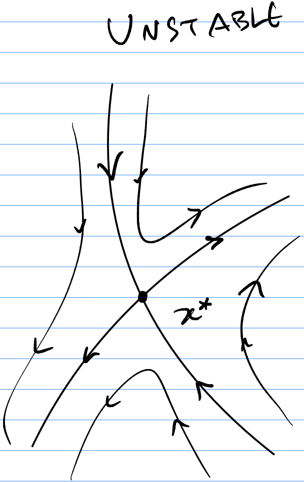
UNSTABLE
x^*

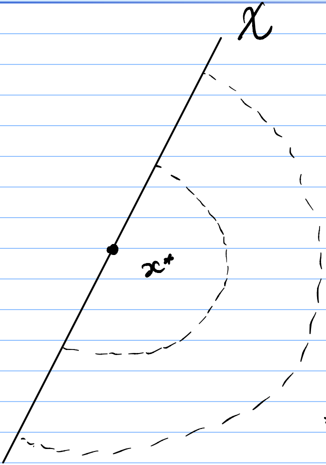

Assume by contradiction that x^* is <u>not</u> Nash

$\Rightarrow \exists\, \alpha \in \operatorname{supp}(x^*)$ $[x^*_\alpha > 0]$

$\exists\, \beta$ s.t. $u_\alpha(x^*) < u_\beta(x^*)$

By cont. of u $\Rightarrow \exists$ some cpt nhd U s.t.

$$\varepsilon + u_\alpha(x) < u_\beta(x) \quad \forall x \in U$$

$\Rightarrow \exists$ IC $x(0)$ s.t. $u_\alpha(x(t)) < u_\beta(x(t)) - \varepsilon$ for all $t \geq 0$

$$\Rightarrow \frac{x_\alpha(t)}{x_\beta(t)} = \exp\left(\int_0^t \underbrace{[u_\alpha(x(s)) - u_\beta(x(s))]}_{\leq -\varepsilon} ds \right) \Rightarrow x_\alpha(t) \to 0 \text{ as } t \to \infty$$

$\Rightarrow\Leftarrow$

Stability

Are all stationary points created equal?

Definition

x^* is **(Lyapunov) stable** if, for every neighborhood U of x^* in $\mathcal{X}$, there exists a neighborhood U' of x^* such that

$$x(0) \in U' \iff x(t) \in U \quad \text{for all } t \geq 0$$

[Trajectories that start close to x^* remain close for all time]

x^* is stable $\implies$ x^* is a Nash equilibrium

Asymptotic stability

Are all Nash equilibria created equal?

Definition

- x^* is **attracting** if $\lim_{t\to\infty} x(t) = x^*$ whenever $x(0)$ is close enough to x^*
- x^* is **asymptotically stable** if it is stable and attracting

Strict Nash (?) = only candidates for asymptotic stability

↳ $\alpha^* = (\alpha_1^*, \dots, \alpha_N^*)$ s.t. $u_{i\alpha_i^*}(x^*) > u_{i\alpha_i}(\alpha_i; x^*_{-i})$

for all $\alpha_i \in \mathcal{A}_i$, $i \in \mathcal{N}$

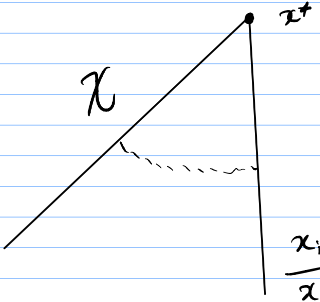

$u_{i\alpha_i^*}(x) > u_{i\alpha_i}(\alpha_i; x_{-i}) + \varepsilon$

for all x close to x^*

$$\frac{x_{i\alpha_i}(t)}{x_{i\alpha_i^*}(t)} = \exp \int_0^t \underbrace{[u_{i\alpha_i}(x(s)) - u_{i\alpha_i^*}(x(s))]}_{\leq -\varepsilon} ds$$

$\leq \exp(-\varepsilon t) \downarrow 0$ as $t \to \infty$

Asymptotic stability

Are all Nash equilibria created equal?

Definition

- x^* is **attracting** if $\lim_{t\to\infty} x(t) = x^*$ whenever $x(0)$ is close enough to x^*
- x^* is **asymptotically stable** if it is stable and attracting

x^* is a **strict** Nash equilibrium $\Longrightarrow$ x^* is asymptotically stable

The "folk theorem" of evolutionary game theory

Theorem (Hofbauer and Sigmund, 2003; Cressman, 2003)

In multi-population random matching games:

- x^* *is a Nash equilibrium* $\implies$ x^* *is stationary*
- x^* *is the limit of an interior trajectory* $\implies$ x^* *is a Nash equilibrium*
- x^* *is stable* $\implies$ x^* *is a Nash equilibrium*
- x^* *is asymptotically stable* $\iff$ x^* *is a strict Nash equilibrium*

The "folk theorem" of evolutionary game theory

Theorem (Hofbauer and Sigmund, 2003; Cressman, 2003)

In single-population random matching games:

- x^* *is a Nash equilibrium* $\implies$ x^* *is stationary*
- x^* *is the limit of an interior trajectory* $\implies$ x^* *is a Nash equilibrium*
- x^* *is stable* $\implies$ x^* *is a Nash equilibrium*
- x^* *is asymptotically stable* $\impliedby$ x^* *is a strict Nash equilibrium*

Attracting full-support equilibria

The replicator dynamics in “good” RPS (win > loss):

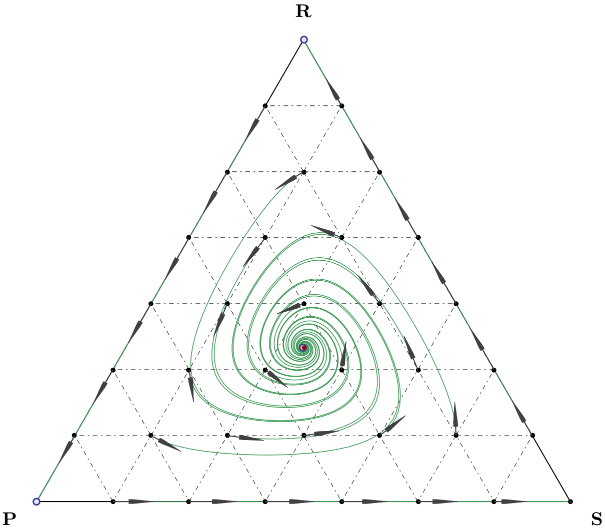

Convergence in potential games

Potential games [Sandholm, 2001]

$$u_{ia_i} = -\frac{\partial V}{\partial x_{ia_i}} \quad \text{for some } \textit{potential function } V\colon \mathcal{X} \to \mathbb{R}$$

NASC (Poincaré):

$$\text{potential} \iff \frac{\partial u_{ia_i}}{\partial x_{ia'_i}} = \frac{\partial u_{ia'_i}}{\partial x_{ia_i}}$$

Positive correlation / Lyapunov property:

$$\frac{dV}{dt} \leq 0 \quad \text{under (RD)}$$

[**Check:** verify this]

Theorem (Sandholm, 2001)

- *In potential games,* (RD) *converges to its set of stationary points*
- *In random matching potential games, interior trajectories of* (RD) *converge to Nash equilibrium*

Non-convergence in zero-sum games

The landscape is very different in zero-sum games:

Non-convergence in zero-sum games

The landscape is very different in zero-sum games:

x^* is full-support equilibrium $\implies$ (RD) admits **constant of motion**

KL divergence: $D_{\mathrm{KL}}(x^*, x) = \sum_i \sum_{a_i} x^*_{ia_i} \log \frac{x^*_{ia_i}}{x_{ia_i}}$

$$D_{KL}(p,x) = \sum_{\alpha : x_\alpha > 0} p_\alpha \log \frac{p_\alpha}{x_\alpha}$$

$$\frac{d}{dt} D_{KL}(p,x) = -\sum_\alpha p_\alpha \frac{\dot{x}_\alpha}{x_\alpha} = \sum_\alpha p_\alpha [u(x) - u_\alpha(x)]$$

$$= \sum_\alpha u_\alpha(x)\,(p_\alpha - x_\alpha)$$

Non-convergence in zero-sum games

The landscape is very different in zero-sum games:

x^* is full-support equilibrium $\implies$ (RD) admits **constant of motion**

KL divergence: $D_{\text{KL}}(x^*, x) = \sum_i \sum_{a_i} x^*_{ia_i} \log \dfrac{x^*_{ia_i}}{x_{ia_i}}$

Theorem (Hofbauer et al., 2009)

Assume a bilinear zero-sum game admits an interior equilibrium. Then:

- *Interior trajectories of* (RD) **do not converge** (*unless stationary*)
- *Time-averages* $\bar{x}(t) = t^{-1} \int_0^t x(\tau)\, d\tau$ **converge to Nash equilibrium**

Convergence of time-averages

The replicator dynamics in a game of Matching Pennies

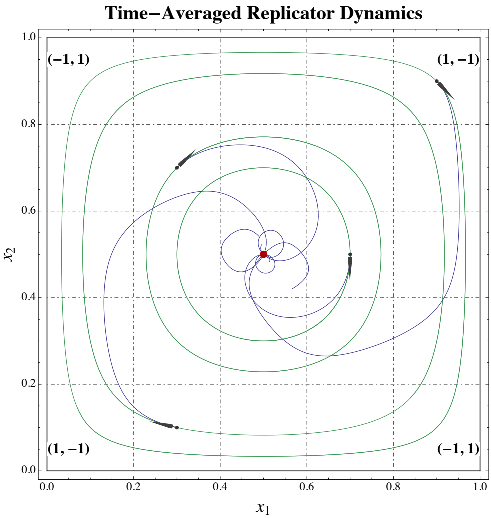

Poincaré recurrence in zero-sum games

Definition (Poincaré)

A dynamical system is **Poincaré recurrent** if almost all solution trajectories return *arbitrarily close* to their starting point *infinitely many times*

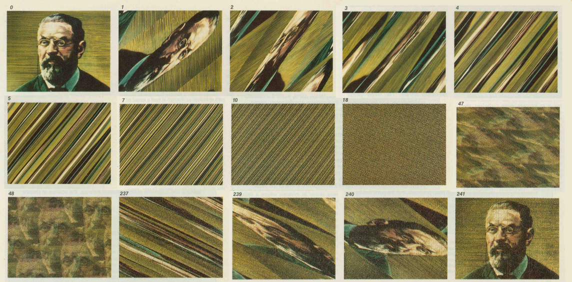

Theorem (Piliouras and Shamma, 2014; M et al., 2018)

(RD) *is recurrent in all bilinear zero-sum games with a full-support equilibrium*

What's missing

- Evolutionary stability [Maynard Smith and Price, 1973]
- Other (classes of) dynamics:
  - Fictitious play [Brown, 1951; Robinson, 1951]
  - Best response dynamics [Gilboa and Matsui, 1991]
  - Imitative / Innovative dynamics [Weibull, 1995; Sandholm, 2010]
  - Higher-order dynamics [Laraki and M, 2013]
- Unexpected / Complex behaviors:
  - Survival of dominated strategies [Hofbauer and Sandholm, 2011]
  - Chaos [Sandholm, 2010]
- Evolution in the presence of uncertainty [Fudenberg and Harris, 1992; Imhof, 2005; M & Moustakas, 2010; M & Viossat, 2016]

References I

G. W. Brown. Iterative solutions of games by fictitious play. In T. C. Coopmans, editor, *Activity Analysis of Productions and Allocation*, 374-376. Wiley, 1951.

S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.

N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

R. Cressman. *Evolutionary Dynamics and Extensive Form Games*. The MIT Press, 2003.

D. Fudenberg and C. Harris. Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57(2):420–441, August 1992.

D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*, volume 2 of *Economic learning and social evolution*. MIT Press, Cambridge, MA, 1998.

D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.

I. Gilboa and A. Matsui. Social stability and equilibrium. *Econometrica*, 59(3):859–867, May 1991.

D. Helbing. A mathematical model for behavioral changes by pair interactions. In G. Haag, U. Mueller, and K. G. Troitzsch, editors, *Economic Evolution and Demographic Change: Formal Models in Social Sciences*, pages 330–348. Springer, Berlin, 1992.

J. Hofbauer and W. H. Sandholm. Survival of dominated strategies under evolutionary dynamics. *Theoretical Economics*, 6(3):341–377, September 2011.

References II

J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK, 1998.

J. Hofbauer and K. Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, July 2003.

J. Hofbauer, S. Sorin, and Y. Viossat. Time average replicator and best reply dynamics. *Mathematics of Operations Research*, 34(2):263–269, May 2009.

L. A. Imhof. The long-run behavior of the stochastic replicator dynamics. *The Annals of Applied Probability*, 15(1B):1019–1045, 2005.

R. Laraki and P. Mertikopoulos. Higher order game dynamics. *Journal of Economic Theory*, 148(6): 2666–2695, November 2013.

T. Lattimore and C. Szepesvári. *Bandit Algorithms*. Cambridge University Press, Cambridge, UK, 2020.

N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108 (2):212–261, 1994.

J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246:15–18, November 1973.

P. Mertikopoulos and A. L. Moustakas. The emergence of rational behavior in the presence of stochastic perturbations. *The Annals of Applied Probability*, 20(4):1359–1388, July 2010.

P. Mertikopoulos and Y. Viossat. Imitation dynamics with payoff shocks. *International Journal of Game Theory*, 45(1-2):291–320, March 2016.

References III

P. Mertikopoulos, C. H. Papadimitriou, and G. Piliouras. Cycles in adversarial regularized learning. In *SODA '18: Proceedings of the 29th annual ACM-SIAM Symposium on Discrete Algorithms*, 2018.

N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

G. Piliouras and J. S. Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via Poincaré recurrence. In *SODA '14: Proceedings of the 25th annual ACM-SIAM Symposium on Discrete Algorithms*, 2014.

J. Robinson. An iterative method for solving a game. *Annals of Mathematics*, 54:296–301, 1951.

A. Rustichini. Optimal properties of stimulus-response learning models. *Games and Economic Behavior*, 29(1-2):244–273, 1999.

L. Samuelson and J. Zhang. Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57: 363–391, 1992.

W. H. Sandholm. Potential games with continuous player sets. *Journal of Economic Theory*, 97:81–108, 2001.

W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA, 2010.

W. H. Sandholm. Population games and deterministic evolutionary dynamics. In H. P. Young and S. Zamir, editors, *Handbook of Game Theory IV*, pages 703–778. Elsevier, 2015.

References IV

K. H. Schlag. Why imitate, and if so, how? a boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78(1):130–156, 1998.

S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

P. D. Taylor and L. B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.

V. G. Vovk. A game of prediction with expert advice. In *COLT '95: Proceedings of the 8th Annual Conference on Computational Learning Theory*, pages 51–60, 1995.

J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.