ΣΤΟΙΧΕΙΑ ΘΕΩΡΙΑΣ ΠΑΙΓΝΙΩΝ ΚΑΙ ΛΗΨΗΣ ΑΠΟΦΑΣΕΩΝ

ΒΑΣΙΚΕΣ ΕΝΝΟΙΕΣ ΚΑΙ ΟΡΙΣΜΟΙ

Παναγιώτης Μερτικόπουλος

Εθνικό και Καποδιστριακό Πανεπιστήμιο Αθηνών

Τμήμα Μαθηματικών

Χειμερινό Εξάμηνο, 2023–2024

Outline

Welcome!

Welcome to SEP19: *Topics in Game Theory*

"The study of rational decision-making"

- **Instructors:** Panayotis Mertikopoulos
- **Meeting times:** Mondays 09:00-13:00
- **e-class:** https://eclass.uoa.gr/courses/MATH806/
- **Sessions:** Focus on general theory with some deep dives / practical sessions (TBD)
- **Grading scheme:** split between end-of-term project (50%) and final (50%)

Course overview

Rough breakdown of the course:

1. **Part 1: Basic elements of game theory**
   - Basic notions: Nash equilibrium, dominated strategies,...
   - Basic notions: Nash equilibrium, dominated strategies,...
   - Game classes: potential games, congestion games, price of anarchy,...
   - Game dynamics: replicator dynamics, exponential weights,...
2. **Part 2: Multi-armed bandits and online optimization**
   - Bandits and regret: regret minimization,...
   - Algorithms: HEDGE, EXP3,...
   - Online convex optimization: regret, convexification,...
   - Algorithms: leader-following policies, gradient / mirror descent,...

Why game theory?

Example 1: A game of roads

A beautiful morning commute in Chicago

The price of congestion

In the US alone, congestion cost $305 billion in 2017 (≈1.6% of GDP)

source: INRIX

- Lost productivity
- Fuel waste
- Environmental impact, quality of life,...

Game of roads

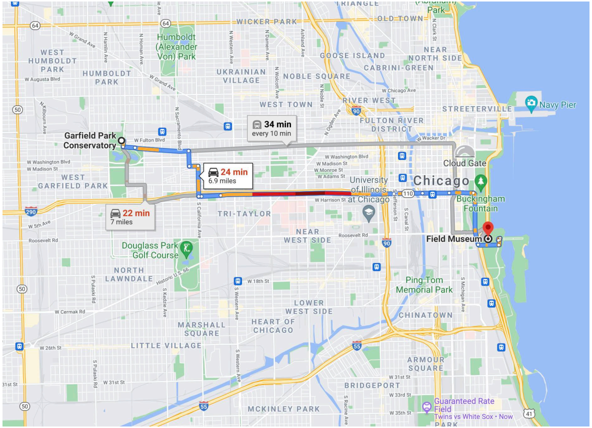

The city of Chicago

- 2,700,000 people
- 1,261,000 daily trips
- 933 nodes
- 2950 edges
- 870,000 o/d pairs
- $\approx 2 * 10^{16}$ paths

A very large game!

Example 2: Spot the fake

Which person is real?

Example 2: Spot the fake

Which person is real?

Spoiler: https://thispersondoesnotexist.com

Neural networks

The workhorse of deep learning:

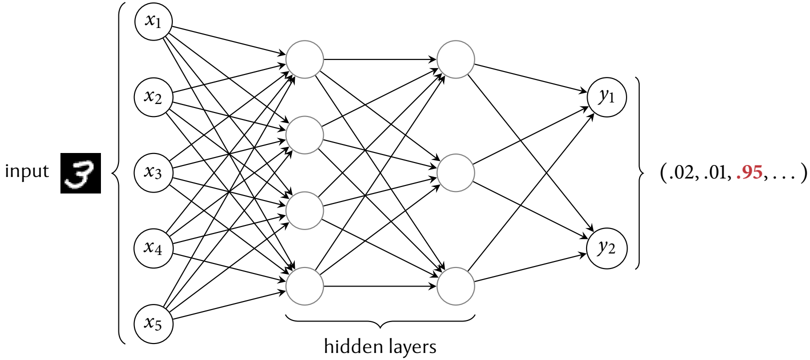

The deep learning revolution: breaking the human perception barrier (2010's)

Neurons

The atoms of any deep learning architecture are its **neurons:**

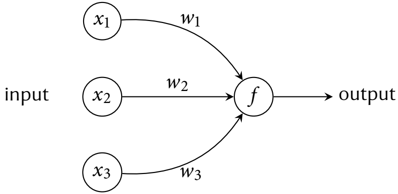

- **Input** could be binary $\{0,1\}$ or real (e.g., average intensity of image)
- Inputs weighed with **weight coefficients** w_i
- Neuron **activates** on value of $f(\sum_i w_i x_i)$

Examples

1. ***Perceptron:*** binary inputs, step function activation
2. ***Sigmoid neuron:*** real inputs, tanh activation
3. ***ReLU:*** real inputs, rectified linear activation ($f(z) = [z]_+$)

The schematics of GANs

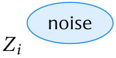

The schematics of GANs

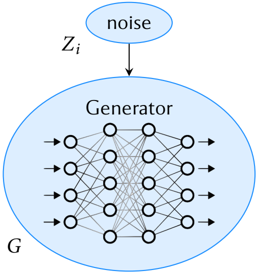

The schematics of GANs

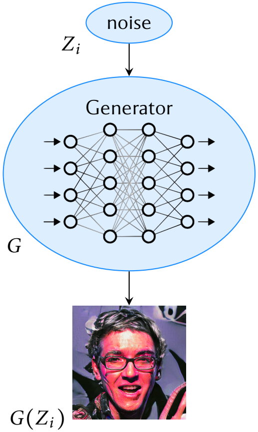

The schematics of GANs

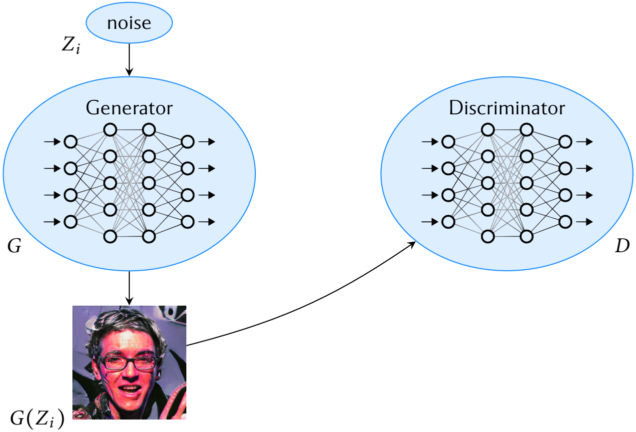

The schematics of GANs

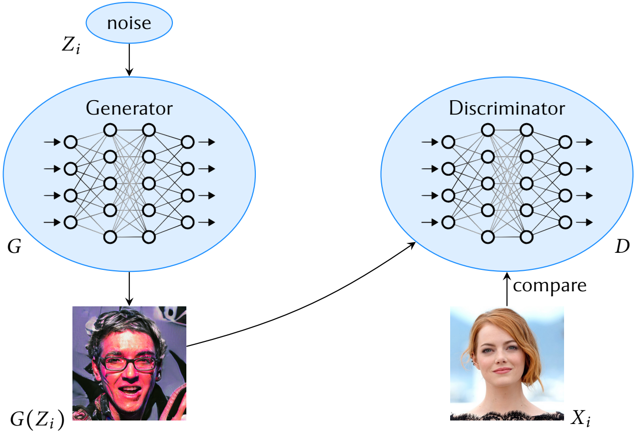

The schematics of GANs

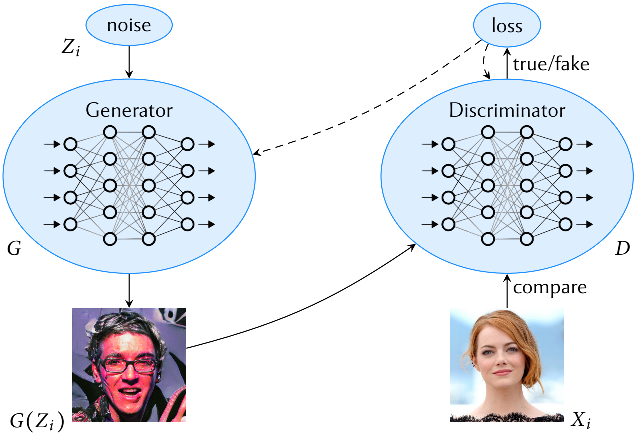

Model likelihood: $\ell(G,D) = \prod_{i=1}^{N} D(X_i) \times \prod_{i=1}^{N} (1 - D(G(Z_i)))$

GAN training

How to find good generators (G) and discriminators (D)?

Discriminator: maximize (log-)likelihood estimation

$$\max_{D \in \mathcal{D}} \log \ell(G, D)$$

Generator: minimize the resulting divergence

$$\min_{G \in \mathcal{G}} \max_{D \in \mathcal{D}} \log \ell(G, D)$$

A (very complex) zero-sum game!

Training landscape

A deep learning loss landscape

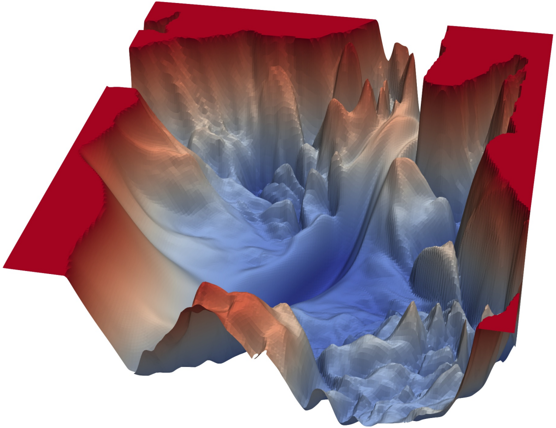

➻ Easier problem: find a needle in a haystack

FailGAN

The game does not always work out:

➼ A StyleGAN after 8 days of training at Nvidia headquarters (!!!)

Questions we'll try to answer

1. **How should we model player interactions?**
   - Urban traffic ≠ transit systems ≠ packet networks ≠ …
   - Rational agents ≠ humans ≠ AI algorithms ≠ …
   - Competition ≠ congestion ≠ coordination ≠ …
2. **What is a desired operational state?**
   - Social optimum ≠ equilibrium ≠ …
   - Static (equilibrium, social optimum) ≠ Bayesian ≠ online (regret) ≠ …
3. **How to compute it?**
   - Calculation ≠ learning ≠ implementation
   - Informational constraints: feedback, bounded rationality, uncertainty, …

Let's play a game

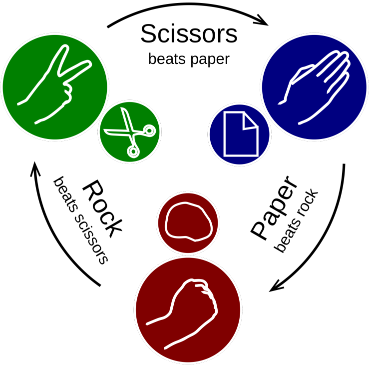

What would you play? How can we model this game mathematically?

Let's play a game, formally

- ***Players:*** "1" and "2"
- ***Actions*** associated to each player: $\mathcal{A}_i = \{R, P, S\}, i = 1, 2$
- ***Payoff matrix*** (win: \$1; lose −\$1; tie \$0):

$$A = \begin{array}{c|ccc} & R & P & S \\ \hline R & 0 & -1 & 1 \\ P & 1 & 0 & -1 \\ S & -1 & 1 & 0 \end{array}$$

- ***Payoff functions:***
  - $u_1 \colon \mathcal{A}_1 \times \mathcal{A}_2 \to \mathbb{R}$ given by $u_1(R, R) = 0$, $u_1(R, P) = -1, \dots$
  - $u_2 \colon \mathcal{A}_1 \times \mathcal{A}_2 \to \mathbb{R}$ given by $u_2(R, R) = 0$, $u_2(R, P) = 1, \dots$

Some basics

What's in a game?

A ***game in normal form*** is a collection of three basic elements:

1. A set of ***players*** $\mathcal{N}$
2. A set of ***actions*** (or ***pure strategies***) $\mathcal{A}_i$ per player $i \in \mathcal{N}$
3. An ensemble of ***payoff functions*** $u_i \colon \mathcal{A} \equiv \prod_j \mathcal{A}_j \to \mathbb{R}$ per player $i \in \mathcal{N}$

Some basics

What's in a game?

A ***game in normal form*** is a collection of three basic elements:

1. A set of ***players*** $\mathcal{N}$
2. A set of ***actions*** (or ***pure strategies***) $\mathcal{A}_i$ per player $i \in \mathcal{N}$
3. An ensemble of ***payoff functions*** $u_i \colon \mathcal{A} \equiv \prod_j \mathcal{A}_j \to \mathbb{R}$ per player $i \in \mathcal{N}$

Important:

- Player set: atomic vs. nonatomic
- Action sets: finite vs. continuous; shared vs. individual; ...

☞ ***NB:*** do not mix game classes!

Taxonomy

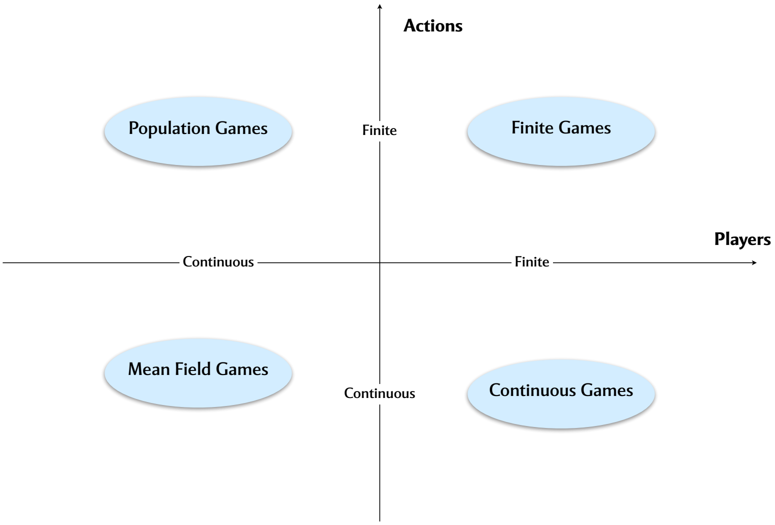
Actions
Population Games
Finite
Finite Games
Players
Continuous
Finite
Mean Field Games
Continuous
Continuous Games

Taxonomy

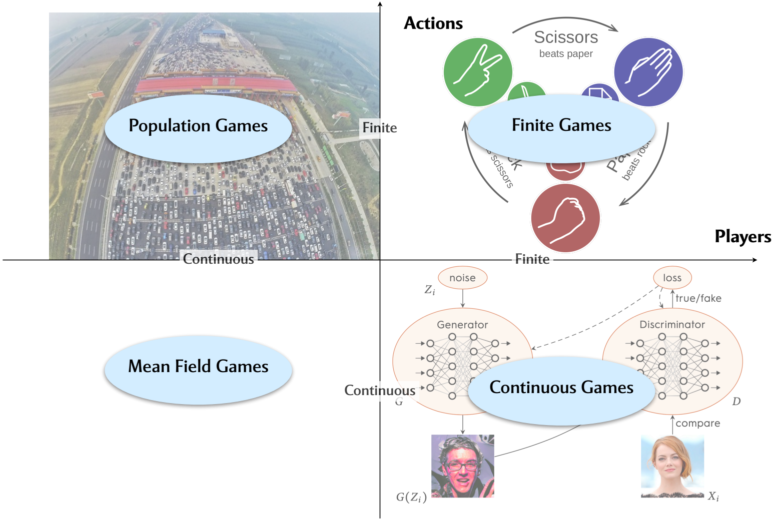
Actions
Players
Finite
Continuous
Finite
Continuous
Population Games
Finite Games
Mean Field Games
Continuous Games
Scissors
beats paper
noise
loss
true/fake
Generator
Discriminator
compare
Z_i
$G(Z_i)$
X_i
D

What's in a game?

Definition (Finite games)

A ***finite game in normal form*** is a collection of the following primitives:

- A finite set of ***players*** $\mathcal{N} = \{1, \dots, N\}$
- A finite set of ***actions*** (or ***pure strategies***) $\mathcal{A}_i$ for each player $i \in \mathcal{N}$
- A ***payoff function*** $u_i \colon \mathcal{A} := \prod_j \mathcal{A}_j \to \mathbb{R}$ for each player $i \in \mathcal{N}$

A game with primitives as above will be denoted as $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$.

Some notes:

- "Normal form" $\rightsquigarrow$ difference with "extensive form" games (Chess, Go,...)
- Handy shorthands: $(a_1, \dots, a_i, \dots a_N) \leftarrow (a_i; a_{-i})$ and $\mathcal{A}_{-i} = \prod_{j \neq i} \mathcal{A}_j$

The Prisoner's Dilemma

Bonnie and Clyde are captured by the authorities and put in separate cells:

- If both betray each other, they both serve 2 years in prison
- If Bonnie betrays but Clyde remains silent, Bonnie goes free and Clyde serves 3 years
- If Bonnie remains silent but Clyde betrays, Bonnie serves 3 years and Clyde goes free
- If neither betrays the other, they both serve 1 year

The Prisoner's Dilemma

Bonnie and Clyde are captured by the authorities and put in separate cells:

- If both betray each other, they both serve 2 years in prison
- If Bonnie betrays but Clyde remains silent, Bonnie goes free and Clyde serves 3 years
- If Bonnie remains silent but Clyde betrays, Bonnie serves 3 years and Clyde goes free
- If neither betrays the other, they both serve 1 year

Normal form representation:

- Players: $\mathcal{N} = \{B, C\}$
- Actions: $\mathcal{A}_B = \mathcal{A}_C = \{\texttt{betray}, \texttt{silent}\}$
- Payoff bimatrix:

$B \downarrow \; C \rightarrow$	betray	silent
betray	$(-2, -2)$	$(0, -3)$
silent	$(-3, 0)$	$(-1, -1)$

Split or steal?

`https://www.youtube.com/watch?v=S0qjK3TWZE8`

- If both players steal, they both get nothing
- If one player steals and the other splits, the one who steals gets everything
- If both players split, they split the prize

Do you `split` or `steal`?

Split or steal?

`https://www.youtube.com/watch?v=S0qjK3TWZE8`

- If both players steal, they both get nothing
- If one player steals and the other splits, the one who steals gets everything
- If both players split, they split the prize

Do you `split` or `steal`?

Normal form representation:

- Players: $\mathcal{N} = \{A, B\}$
- Actions: $\mathcal{A}_A = \mathcal{A}_B = \{\texttt{split}, \texttt{steal}\}$
- Payoff bimatrix:

$A \downarrow\ B \rightarrow$	split	steal
split	$(\$6800, \$6800)$	$(0, \$13600)$
steal	$(\$13600, 0)$	$(0, 0)$

The battle of the sexes

Robin and Charlie want to go out for the evening:

- Robin prefers to go to a movie
- Charlie prefers to go to the theater
- They both prefer being together instead of alone

The battle of the sexes

Robin and Charlie want to go out for the evening:

- Robin prefers to go to a movie
- Charlie prefers to go to the theater
- They both prefer being together instead of alone

Normal form representation:

- Players: $\mathcal{N} = \{R, C\}$
- Actions: $\mathcal{A}_R = \mathcal{A}_C = \{\texttt{movie}, \texttt{theater}\}$
- Payoff bimatrix:

$R \downarrow \ C \rightarrow$	movie	theater
movie	$(3, 2)$	$(0, 0)$
theater	$(0, 0)$	$(2, 3)$

The collision game

Robin and Charlie arrive at an uncontrolled intersection:

- If they both drive through, they crash
- If they both yield, they may wait forever
- If one yields and the other drives through, the latter loses less time

The collision game

Robin and Charlie arrive at an uncontrolled intersection:

- If they both drive through, they crash
- If they both yield, they may wait forever
- If one yields and the other drives through, the latter loses less time

Normal form representation:

- Players: $\mathcal{N} = \{R, C\}$
- Actions: $\mathcal{A}_R = \mathcal{A}_C = \{\texttt{drive}, \texttt{yield}\}$
- Payoff bimatrix:

$R \downarrow\ C \rightarrow$	drive	yield
drive	$(-100, -100)$	$(2, 1)$
yield	$(1, 2)$	$(0, 0)$

Outline

Dominated strategies

Sometimes, an action may yield consistently suboptimal payoffs

Definition (Dominated strategies)

1. A strategy $a_i \in \mathcal{A}_i$ is ***strictly dominated*** by $a_i' \in \mathcal{A}_i$ if

$$u_i(a_i; a_{-i}) < u_i(a_i'; a_{-i}) \quad \text{for all } a_{-i} \in \mathcal{A}_{-i}$$

2. A strategy $a_i \in \mathcal{A}_i$ is ***weakly dominated*** by $a_i' \in \mathcal{A}_i$ if

$$u_i(a_i; a_{-i}) \leq u_i(a_i'; a_{-i}) \quad \text{for all } a_{-i} \in \mathcal{A}_{-i}$$

and $u_i(a_i; a_{-i}) < u_i(a_i'; a_{-i})$ for some $a_{-i} \in \mathcal{A}_{-i}$.

Notation:

- a_i is *strictly* dominated by a_i': $a_i \prec a_i'$
- a_i is *weakly* dominated by a_i': $a_i \precsim a_i'$

Examples, revisited

The prisoner's dilemma:

$R\downarrow\ C\rightarrow$	betray	silent
betray	$(-2,-2)$	$(0,-3)$
silent	$(-3,0)$	$(-1,-1)$

Split or steal:

$R\downarrow\ C\rightarrow$	split	steal
split	$(\$6800,\$6800)$	$(0,\$13600)$
steal	$(\$13600,0)$	$(0,0)$

Battle of the sexes:

$R\downarrow\ C\rightarrow$	movie	theater
movie	$(3,2)$	$(0,0)$
theater	$(0,0)$	$(2,3)$

Iteratively dominated strategies

A larger game:

$$\begin{matrix} (9,4) & (5,3) & (3,2) \\ (0,1) & (4,6) & (6,0) \\ (2,1) & (3,5) & (2,4) \end{matrix}$$

Iteratively dominated strategies

A larger game:

$$\begin{matrix} (9,4) & (5,3) & (3,2) \\ (0,1) & (4,6) & (6,0) \\ (2,1) & (3,5) & (2,4) \end{matrix}$$

Definition

1. A strategy is called ***iteratively dominated*** if it becomes dominated after successive elimination of dominated strategies.
2. A game is called ***dominance-solvable*** if the successive elimination of dominated strategies leads to a singleton.

Best responses

What if only the strategy of the opposing player(s) is known?

Definition (Best responses)

The strategy $a_i^* \in \mathcal{A}_i$ is a ***best response*** to $a_{-i} \in \mathcal{A}_{-i}$ if

$$u_i(a_i^*; a_{-i}) \geq u_i(a_i; a_{-i}) \quad \text{for all } a_i \in \mathcal{A}_i$$

or, equivalently, if

$$a_i^* \in \arg\max_{a_i \in \mathcal{A}_i} u_i(a_i; a_{-i}).$$

The set-valued function $\mathrm{BR}_i \colon \mathcal{A}_{-i} \rightrightarrows \mathcal{A}_i$ given by

$$\mathrm{BR}_i(a_{-i}) = \arg\max_{a_i \in \mathcal{A}_i} u_i(a_i; a_{-i})$$

is called the ***best-response correspondence***.

Examples

The prisoner's dilemma:

$R\downarrow\ C\rightarrow$	betray	silent
betray	$(-2,-2)$	$(0,-3)$
silent	$(-3,0)$	$(-1,-1)$

Split or steal:

$R\downarrow\ C\rightarrow$	split	steal
split	$(\$6800,\$6800)$	$(0,\$13600)$
steal	$(\$13600,0)$	$(0,0)$

Battle of the sexes:

$R\downarrow\ C\rightarrow$	movie	theater
movie	$(3,2)$	$(0,0)$
theater	$(0,0)$	$(2,3)$

Dominated strategies and best responses

Some more examples of best responses

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

Dominated strategies and best responses

Some more examples of best responses

$$\begin{matrix} (9,4) & (5,3) & (3,2) \\ (0,1) & (4,6) & (6,0) \\ (2,1) & (3,5) & (2,8) \end{matrix}$$

Best responses cannot contain dominated strategies

Dominated strategies and best responses

Some more examples of best responses

$$\begin{matrix} (9,4) & (5,3) & (3,2) \\ (0,1) & (4,6) & (6,0) \\ (2,1) & (3,5) & (2,8) \end{matrix}$$

Best responses cannot contain dominated strategies

➪ What about *weakly* dominated strategies?

Nash equilibrium

Equilibrium: best-responding to each other's actions

Definition (Nash equilibrium)

An action profile $a^* = (a_1^*, \dots, a_N^*)$ is a ***Nash equilibrium*** if

$$a_i^* \in \mathrm{BR}_i(a_{-i}^*) \quad \text{for all } i \in \mathcal{N}$$

or, equivalently, if

$$u_i(a_i^*; a_{-i}^*) \geq u_i(a_i; a_{-i}^*) \quad \text{for all } a_i \in \mathcal{A}_i \text{ and all } i \in \mathcal{N}.$$

Intuition:

- ***Stability:*** no player has an incentive to deviate
- ***Unilateral resilience:*** stable against *individual* player deviations, not multi-player ones

Examples, revisited

The prisoner's dilemma:

$R \downarrow C \rightarrow$	betray	silent
betray	$(-2,-2)$	$(0,-3)$
silent	$(-3,0)$	$(-1,-1)$

Split or steal:

$R \downarrow C \rightarrow$	split	steal
split	$(\$6800,\$6800)$	$(0,\$13600)$
steal	$(\$13600,0)$	$(0,0)$

Battle of the sexes:

$R \downarrow C \rightarrow$	movie	theater
movie	$(3,2)$	$(0,0)$
theater	$(0,0)$	$(2,3)$

RPS, revisited

How about Rock-Paper-Scissors?

	R	P	S
R	0	−1	1
P	1	0	−1
S	−1	1	0

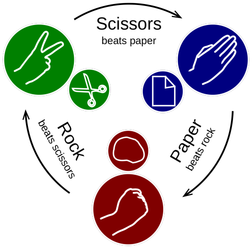

RPS, revisited

How about Rock-Paper-Scissors?

	R	P	S
R	0	–1	1
P	1	0	–1
S	–1	1	0

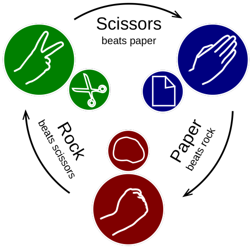

Nash equilibria don't always exist!

Outline

Mixed strategies

Instead of playing pure strategies, players could ***mix*** their actions:

- ***Mixed strategy*** of player $i \in \mathcal{N}$: probability distribution x_i on $\mathcal{A}_i$
- ***Notation:*** x_{ia_i} = prob. that player i selects $a_i \in \mathcal{A}_i$
- ***Strategy space*** of player i:

$$\mathcal{X}_i := \Delta(\mathcal{A}_i) = \left\{x_i \in \mathbb{R}^{\mathcal{A}_i} : x_{ia_i} \geq 0 \text{ and } \sum_{a_i \in \mathcal{A}_i} x_{ia_i} = 1\right\}$$

➻ $\Delta(\mathcal{A}_i) \rightsquigarrow$ simplex spanned by $\mathcal{A}_i$

- ***Support*** of x_i: actions that are played with positive probability under x_i

$$\text{supp}(x_i) := \{a_i \in \mathcal{A}_i : x_{ia_i} > 0\}$$

- x_i is ***pure*** when $\text{supp}(x_i)$ is a singleton, i.e.,

$$\text{supp}(x_i) = \{a_i\} \quad \text{for some } a_i \in \mathcal{A}_i$$

➻ Origin of the term "pure strategies"

RPS, revisited

Playing with mixed strategies:

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

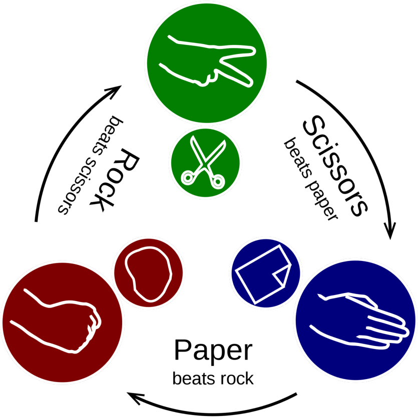

RPS, revisited

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$

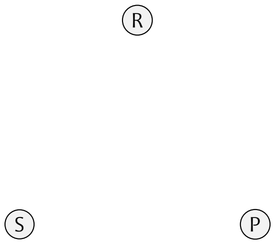

RPS, revisited

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{R, P, S\}$
- Mixed strategy space: $\mathcal{X}_i = \Delta\{R, P, S\}$

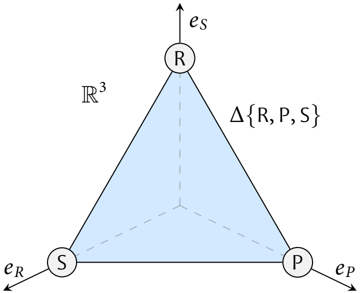

RPS, revisited

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$
- Mixed strategy space: $\mathcal{X}_i = \Delta\{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$
- Choose mixed strategy $x_i \in \mathcal{X}_i$

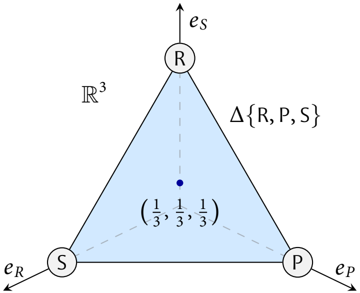

RPS, revisited

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{R, P, S\}$
- Mixed strategy space: $\mathcal{X}_i = \Delta\{R, P, S\}$
- Choose mixed strategy $x_i \in \mathcal{X}_i$
- Choose action $a_i \sim x_i$

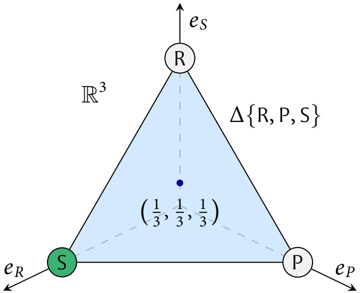

Mixed strategies (collective)

When all players mix their actions:

- Each player $i \in \mathcal{N}$ uses a mixed strategy $x_i \in \mathcal{X}_i$
- Prob. of selecting the action profile $a = (a_1, \dots, a_N) \in \mathcal{A} = \prod_j \mathcal{A}_j$:

$$x_{a_1,\dots,a_N} = \prod_{j\in\mathcal{N}} x_{ja_j}$$

- Prob. of selecting $a_{-i} \in \mathcal{A}_{-i}$:

$$x_{-i;a_{-i}} = \prod_{j\neq i} x_{ja_j}$$

Mixed strategies (collective)

When all players mix their actions:

- Each player $i \in \mathcal{N}$ uses a mixed strategy $x_i \in \mathcal{X}_i$
- Prob. of selecting the action profile $a = (a_1, \dots, a_N) \in \mathcal{A} = \prod_j \mathcal{A}_j$:

$$x_{a_1,\dots,a_N} = \prod_{j\in\mathcal{N}} x_{ja_j}$$

- Prob. of selecting $a_{-i} \in \mathcal{A}_{-i}$:

$$x_{-i;a_{-i}} = \prod_{j\neq i} x_{ja_j}$$

- ***Mixed strategy profile:***

$$x = (x_1, \dots, x_N) \in \mathcal{X} := \prod_{i\in\mathcal{N}} \mathcal{X}_i$$

- ***Mixed strategy profile of i's opponents:***

$$x_{-i} = (x_1, \dots, \cancel{x_i}, \dots, x_N) \in \mathcal{X}_{-i} := \prod_{j\neq i} \mathcal{X}_j$$

☞ **NB:** $\mathcal{X} = \prod_j \Delta(\mathcal{A}_j) \neq \Delta(\prod_j \mathcal{A}_j) = \Delta(\mathcal{A})$

➼ *mixed* vs. *correlated* strategies

Expected payoffs

Expected payoffs under mixed strategies:

- ***Expected payoff to a player*** under a mixed strategy profile:

$$u_i(x) = \sum_{a_1 \in \mathcal{A}_1} \cdots \sum_{a_N \in \mathcal{A}_N} x_{1,a_1} \cdots x_{N,a_N} \, u_i(a_1, \dots, a_N)$$

  or, in terms of other players' strategies:

$$u_i(x_i; x_{-i}) = \sum_{a_i \in \mathcal{A}_i} \sum_{a_{-i} \in \mathcal{A}_{-i}} x_{ia_i} x_{-i;a_{-i}} \, u_i(a_i; a_{-i})$$

- ***Expected payoff to a pure strategy*** under a mixed strategy profile:

$$v_{ia_i}(x) := u_i(a_i; x_{-i}) = \sum_{a_{-i} \in \mathcal{A}_{-i}} x_{-i;a_{-i}} u_i(a_i; a_{-i})$$

Expected payoffs

Expected payoffs under mixed strategies:

- ***Expected payoff to a player*** under a mixed strategy profile:

$$u_i(x) = \sum_{a_1 \in \mathcal{A}_1} \cdots \sum_{a_N \in \mathcal{A}_N} x_{1,a_1} \cdots x_{N,a_N} \, u_i(a_1, \dots, a_N)$$

  or, in terms of other players' strategies:

$$u_i(x_i; x_{-i}) = \sum_{a_i \in \mathcal{A}_i} \sum_{a_{-i} \in \mathcal{A}_{-i}} x_{i a_i} x_{-i; a_{-i}} \, u_i(a_i; a_{-i})$$

- ***Expected payoff to a pure strategy*** under a mixed strategy profile:

$$v_{i a_i}(x) := u_i(a_i; x_{-i}) = \sum_{a_{-i} \in \mathcal{A}_{-i}} x_{-i; a_{-i}} u_i(a_i; a_{-i})$$

- ***Mixed payoff vectors:***

$$v_i(x) = (v_{i a_i}(x))_{a_i \in \mathcal{A}_i} = (u_i(a_i; x_{-i}))_{a_i \in \mathcal{A}_i}$$

  so

$$u_i(x) = \langle v_i(x), x_i \rangle$$

☞ **NB:** u_i is **linear** in x_i; $v_{i a_i}$ and v_i are **independent** of x_i

Go-to example: Rock-Paper-Scissors

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{R, P, S\}$
- Mixed strategies: $x_i \in \mathcal{X}_i$

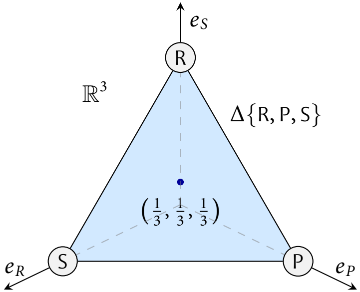

Go-to example: Rock-Paper-Scissors

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$
- Mixed strategies: $x_i \in \mathcal{X}_i$

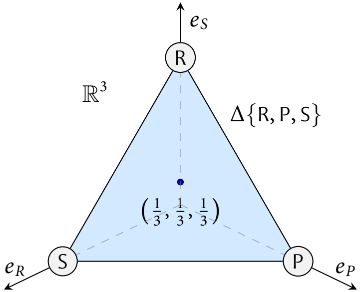

Mixed strategy payoffs:

$$\begin{aligned}
u_1(x_1, x_2) &= x_{1,\mathrm{R}} x_{2,\mathrm{R}} \cdot (0) + x_{1,\mathrm{R}} x_{2,\mathrm{P}} \cdot (-1) + x_{1,\mathrm{R}} x_{2,\mathrm{S}} \cdot (1) \\
&\quad + x_{1,\mathrm{P}} x_{2,\mathrm{R}} \cdot (1) + x_{1,\mathrm{P}} x_{2,\mathrm{P}} \cdot (0) + x_{1,\mathrm{P}} x_{2,\mathrm{S}} \cdot (-1) \\
&\quad + x_{1,\mathrm{S}} x_{2,\mathrm{R}} \cdot (-1) + x_{1,\mathrm{S}} x_{2,\mathrm{P}} \cdot (1) + x_{1,\mathrm{S}} x_{2,\mathrm{S}} \cdot (0) \\
&= x_{1,\mathrm{R}}(x_{2,\mathrm{S}} - x_{2,\mathrm{P}}) + x_{1,\mathrm{P}}(x_{2,\mathrm{R}} - x_{2,\mathrm{S}}) + x_{1,\mathrm{S}}(x_{2,\mathrm{P}} - x_{2,\mathrm{R}}) \\
&= x_1^\top A x_2 \\
u_2(x_1, x_2) &= -u_1(x_1, x_2)
\end{aligned}$$

Mixed extensions

Definition (Mixed extension of a finite game)

The ***mixed extension*** of a finite game $\Gamma = \Gamma(\mathcal{N}, \mathcal{A}, u)$ is the ***continuous*** game $\Delta(\Gamma)$ with

- Players $i \in \mathcal{N} = \{1, \dots, N\}$
- Actions $x_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i)$ per player $i \in \mathcal{N}$
- Payoff functions $u_i : \mathcal{X} \to \mathbb{R}$, $i \in \mathcal{N}$

Notes:

- ***Continuous game:*** game with *continuous* action spaces (here $\mathcal{X}_i$ instead of $\mathcal{A}_i$)
- ***Context:*** when clear, we will not distinguish between Γ and $\Delta(\Gamma)$

Mixed best responses

Extending the notion of best-responding to mixed strategies

Definition (Mixed best responses)

The mixed strategy $x_i^* \in \mathcal{X}_i$ is a ***best response*** to the mixed profile $x_{-i} \in \mathcal{X}_{-i}$ if

$$u_i(x_i^*; x_{-i}) \geq u_i(x_i; x_{-i}) \quad \text{for all } x_i \in \mathcal{X}_i$$

or, equivalently, if

$$x_i^* \in \arg\max_{x_i \in \mathcal{X}_i} u_i(x_i; x_{-i}) = \arg\max_{x_i \in \mathcal{X}_i} \langle v_i(x), x_i \rangle$$

As before, we write $\mathrm{BR}_i(x_{-i}) = \arg\max_{x_i \in \mathcal{X}_i} u_i(x_i; x_{-i})$.

Notes:

- ***Structure:*** $\mathrm{BR}_i(x_{-i})$ is always a face of $\mathcal{X}_i$ Why?
- ***Notation:*** rely on context to distinguish between pure / mixed best responses

Outline

Go-to example: Rock-Paper-Scissors

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{R, P, S\}$
- Mixed strategies: $x_i^* \in \mathcal{X}_i$

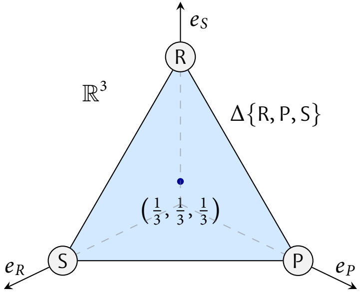

Go-to example: Rock-Paper-Scissors

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$
- Mixed strategies: $x_i^* \in \mathcal{X}_i$

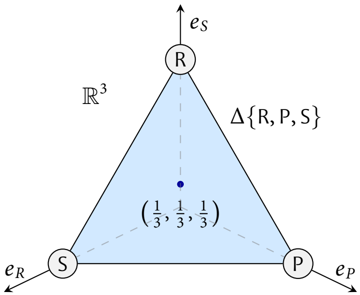

Mixed strategy payoffs when $x_1^* = x_2^* = (1/3, 1/3, 1/3)$:

$$u_1(x_1^*, x_2^*) = \frac{1}{3}\left(\frac{1}{3} - \frac{1}{3}\right) + \frac{1}{3}\left(\frac{1}{3} - \frac{1}{3}\right) + \frac{1}{3}\left(\frac{1}{3} - \frac{1}{3}\right) = 0 = u_2(x_1^*, x_2^*)$$

Go-to example: Rock-Paper-Scissors

Playing with mixed strategies:

- Players: $\mathcal{N} = \{1, 2\}$
- Actions: $\mathcal{A}_i = \{\mathrm{R}, \mathrm{P}, \mathrm{S}\}$
- Mixed strategies: $x_i^* \in \mathcal{X}_i$

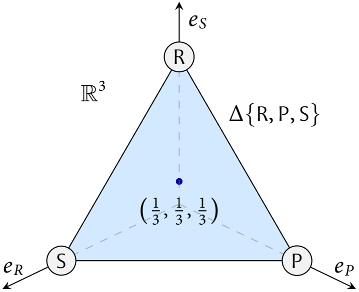

Mixed strategy payoffs when $x_1^* = x_2^* = (1/3, 1/3, 1/3)$:

$$u_1(x_1^*, x_2^*) = \frac{1}{3}\left(\frac{1}{3} - \frac{1}{3}\right) + \frac{1}{3}\left(\frac{1}{3} - \frac{1}{3}\right) + \frac{1}{3}\left(\frac{1}{3} - \frac{1}{3}\right) = 0 = u_2(x_1^*, x_2^*)$$

In fact:

$$u_1(x_1, x_2^*) = 0 = u_2(x_1^*, x_2) \quad \text{for all } x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2$$

so

$$x_1^* \in \mathrm{BR}_1(x_2^*) \quad \text{and} \quad x_2^* \in \mathrm{BR}_2(x_1^*)$$

Nash equilibrium in mixed strategies

Extending the notion of equilibrium to mixed strategies

Definition (Nash equilibrium)

A strategy profile $x^* = (x_1^*, \dots, x_N^*)$ is a ***Nash equilibrium*** if

$$x_i^* \in \mathrm{BR}_i(x_{-i}^*) \quad \text{for all } i \in \mathcal{N}$$

or, equivalently, if

$$u_i(x_i^*; x_{-i}^*) \geq u_i(x_i; x_{-i}^*) \quad \text{for all } x_i \in \mathcal{X}_i \text{ and all } i \in \mathcal{N}.$$

Notes:

- ***Unilateral stability:*** ceteris paribus, no player has an incentive to deviate
- If x^* is pure $\implies$ ***pure Nash equilibrium*** — otherwise "*mixed*"
- If ">" instead of "≥" for $x_i \neq x_i^*$ $\implies$ ***strict Nash equilibrium***

☞ **Prove:** x^* is strict $\iff$ $\mathrm{BR}_i(x_{-i}^*)$ is a singleton for all $i \in \mathcal{N}$

Nash's theorem

RPS admits a Nash equilibrium in mixed strategies - is this always the case?

Nash's theorem

RPS admits a Nash equilibrium in mixed strategies - is this always the case?

Theorem (Nash, 1950)

Every finite game admits a Nash equilibrium in mixed strategies.

Notes:

- ***Support:*** Nash's theorem **does not** specify the support or other properties
- ***Oddness:*** generically odd number of equilibria ➬ Wilson (1971)
- ***Index:*** generically, if m *pure* equilibria, at least $m-1$ *mixed* equilibria ➬ Ritzberger (1994)

Proof, Part I

Skeleton of the proof:

- Introduce collective best-response correspondence $\mathrm{BR}\colon \mathcal{X} \rightrightarrows \mathcal{X}$ given by

$$\mathrm{BR}(x) = (\mathrm{BR}_i(x_{-i}))_{i=1,\dots,N}$$

- x^* is a Nash equilibrium $\iff x^* \in \mathrm{BR}(x^*)$

Proof, Part I

Skeleton of the proof:

- Introduce collective best-response correspondence $\mathrm{BR}\colon \mathcal{X} \rightrightarrows \mathcal{X}$ given by

$$\mathrm{BR}(x) = (\mathrm{BR}_i(x_{-i}))_{i=1,\dots,N}$$

- x^* is a Nash equilibrium $\iff x^* \in \mathrm{BR}(x^*)$
- Invoke Kakutani's fixed-point theorem for set-valued functions.

Theorem (Kakutani, 1941)

Let $\mathcal{C}$ be a nonempty compact convex subset of $\mathbb{R}^d$, and let $F\colon \mathcal{C} \rightrightarrows \mathcal{C}$ be a set-valued function such that:

(P1) $F(x)$ *is nonempty, closed and convex for all* $x \in \mathcal{C}$

(P2) F *is* **upper hemicontinuous** *at all* $x \in \mathcal{C}$, *i.e.,* $\tilde{x} \in F(x)$ *whenever* $x_t \to x$ *and* $\tilde{x}_t \to \tilde{x}$ *for sequences* $x_t \in \mathcal{C}$ *and* $\tilde{x}_t \in F(x_t)$.

Then there exists some $x^* \in \mathcal{C}$ *such that* $x^* \in F(x^*)$.

➡ Upper hemicontinuity ⇝ closed graph

Proof, Part II

Verify the conditions of Kakutani's theorem for $\mathcal{C} \leftarrow \mathcal{X}$ and $F \leftarrow \mathrm{BR}$:

(P1) $\mathrm{BR}(x)$ is a face of $\mathcal{X}$, so it is nonempty, closed and convex [Why?]

(P2) Argue by contradiction

- Suppose there exist sequences $x_t, \tilde{x}_t \in \mathcal{X}$, $t = 1, 2, \dots$, such that $x_t \to x$, $\tilde{x}_t \to \tilde{x}$ and $\tilde{x}_t \in \mathrm{BR}(x_t)$, but $\tilde{x} \notin \mathrm{BR}(x)$.
- Then there exists a player $i \in \mathcal{N}$ and a deviation $x'_i \in \mathcal{X}_i$ such that

$$u_i(x'_i; x_{-i}) > u_i(\tilde{x}_i; x_{-i})$$

- But since $\tilde{x}_{i,t} \in \mathrm{BR}(x_{-i,t})$ by assumption, we also have:

$$u_i(x'_i; x_{-i,t}) \leq u_i(\tilde{x}_{i,t}; x_{-i,t})$$

- Since $x_t \to x$, $\tilde{x}_t \to \tilde{x}$ and u_i is continuous, taking limits gives

$$u_i(x'_i; x_{-i}) \leq u_i(\tilde{x}_i; x_{-i})$$

which contradicts our original assumption. □

Outline

Potential games and best responses

Going back to pure strategies:

- *In single-player games:* Nash equilibria (maximizers) trivially exist
- *In multi-player games:* not true

Bridge between single- and multi-player settings?

Potential games and best responses

Going back to pure strategies:

- ***In single-player games:*** Nash equilibria (maximizers) trivially exist
- ***In multi-player games:*** not true

Bridge between single- and multi-player settings?

Definition (Potential games; Monderer & Shapley, 1996)

A finite game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ is a ***potential game*** if there exists a function $\Phi\colon \mathcal{A} \to \mathbb{R}$ such that

$$u_i(a_i'; a_{-i}) - u_i(a_i; a_{-i}) = \Phi(a_i'; a_{-i}) - \Phi(a_i; a_{-i})$$

for all $a, a' \in \mathcal{A}$ and all $i \in \mathcal{N}$.

Potential games and best responses

Going back to pure strategies:

- *In single-player games:* Nash equilibria (maximizers) trivially exist
- *In multi-player games:* not true

Bridge between single- and multi-player settings?

Definition (Potential games; Monderer & Shapley, 1996)

A finite game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ is a ***potential game*** if there exists a function $\Phi\colon \mathcal{A} \to \mathbb{R}$ such that

$$u_i(a_i'; a_{-i}) - u_i(a_i; a_{-i}) = \Phi(a_i'; a_{-i}) - \Phi(a_i; a_{-i})$$

for all $a, a' \in \mathcal{A}$ and all $i \in \mathcal{N}$.

Examples

- Battle of the sexes
- Congestion games (more later...)

Basic properties

Existence of equilibria:

- Any ***global maximizer*** $a^* \in \arg\max \Phi$ of Φ is a pure Nash equilibrium

Basic properties

Existence of equilibria:

- Any ***global maximizer*** $a^* \in \arg\max \Phi$ of Φ is a pure Nash equilibrium
- Any ***unilateral maximizer*** $a^* \in \mathcal{A}$ of Φ is a pure Nash equilibrium
- ***Unilateral maximizers:***

$$\Phi(a^*) \geq \Phi(a_i; a^*_{-i}) \quad \text{for all } a_i \in \mathcal{A}_i \text{ and all } i \in \mathcal{N}$$

Basic properties

Existence of equilibria:

- Any ***global maximizer*** $a^* \in \arg\max \Phi$ of Φ is a pure Nash equilibrium
- Any ***unilateral maximizer*** $a^* \in \mathcal{A}$ of Φ is a pure Nash equilibrium
- ***Unilateral maximizers:***

$$\Phi(a^*) \geq \Phi(a_i; a^*_{-i}) \quad \text{for all } a_i \in \mathcal{A}_i \text{ and all } i \in \mathcal{N}$$

When is a game a potential one?

Proposition

Γ is a potential game if and only if

$$\nabla_{x_j} v_i(x) = \nabla_{x_i} v_j(x) \quad \textit{for all } x \in \mathcal{X} \textit{ and all } i, j \in \mathcal{N}$$

where $v_i(x) = (u_i(a_i; x_{-i}))_{a_i \in \mathcal{A}_i}$ is the mixed payoff vector of player $i \in \mathcal{N}$.

Best-response dynamics

A natural updating process:

- Players may choose a new action at each $t = 1, 2, \ldots$
- Players best-respond if this ***strictly*** increases their payoff

Definition (Best-response dynamics)

The ***best-response dynamics*** are defined by the recursion

$$a_{i_t,t+1} \begin{cases} \in \mathrm{BR}_{i_t}(a_{-i_t,t}) & \text{if } a_{i_t,t} \notin \mathrm{BR}_{i_t}(a_{-i_t,t}) \\ = a_{i_t,t} & \text{otherwise} \end{cases} \tag{BRD}$$

where i_t is any player that updates at stage t.

Notes:

- ***Simultaneous:*** all players update simultaneously
- ***Iterative:*** players update in a round robin fashion
- ***Randomized:*** random subset of players updates at any given stage

Convergence

Does (BRD) converge?

Convergence

Does (BRD) converge?

✗ **No - and different modes of updating don't help** ➻ Think RPS

Convergence

Does (BRD) converge?

✗ No - and different modes of updating don't help

➻ Think RPS

But good convergence properties in potential games:

Proposition (Monderer & Shapley, 1996)

Let Γ be a finite potential game. Then the iterative version of (BRD) *converges to a pure Nash equilibrium after finitely many steps.*

Convergence

Does (BRD) converge?

✗ **No - and different modes of updating don't help** ➾ Think RPS

But good convergence properties in potential games:

Proposition (Monderer & Shapley, 1996)

Let Γ be a finite potential game. Then the iterative version of (BRD) converges to a pure Nash equilibrium after finitely many steps.

Notes:

- ***Simple proof:*** potential before and after an update is

$$\Phi(a_i^+; a_{-i}) - \Phi(a_i; a_{-i}) = u_i(a_i^+; a_{-i}) - u_i(a_i; a_{-i}) > 0$$

  whenever $a_i^+ \neq a_i \implies$ no action profile is visited twice $\implies$ the process stops
- ***Iterative vs. simultaneous:*** the distinction matters, **simultaneous** (BRD) **may cycle**

References I

[1] Kakutani, S. A generalization of Brouwer's fixed point theorem. *Duke Mathematical Journal*, 8(3):457-459, September 1941.

[2] Monderer, D. and Shapley, L. S. Potential games. *Games and Economic Behavior*, 14(1):124 - 143, 1996.

[3] Nash, J. F. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the USA*, 36:48-49, 1950.

[4] Ritzberger, K. The theory of normal form games from the differentiable viewpoint. *International Journal of Game Theory*, 23:207-236, September 1994.

[5] Wilson, R. Computing equilibria of n-person games. *SIAM Journal on Applied Mathematics*, 21:80-87, 1971.