Introduction to Game Theory

Isabel Amigo

IMT Atlantique

October 2017

Notes:

Game Theory in a Nutshell

Game theory: analytical tools designed to help understand the phenomena observed when decision-makers interact (see e.g. [3])

Notes:

WiFi vs 3G access

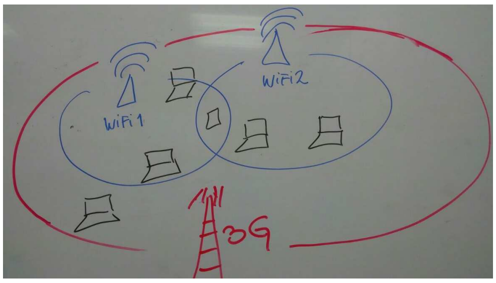

Notes:

- Lets see some examples
- Imagine a simplified situation where users have different options as access technology
- For instance, the fact of choosing one or other network by each user, has consequences in the quality experienced by all users
- Decisions are however taken individually by each user
- What will occur?

TCP congestion control

Correct vs Malicious implementation

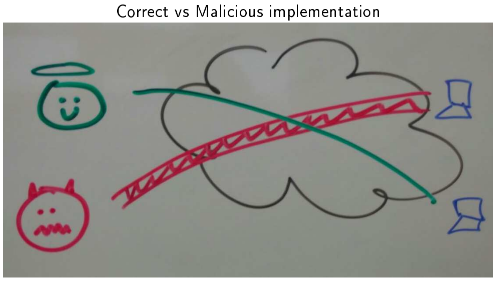

Notes:

- Lets imagine the 2-player version of TCP
- Quick refreshment: TCP has a control algorithmic for congestion avoidance, which is based on the detection of a loss. A loss is interpreted as a sign of congestion, and thus the sending window is diminished (backoff).
- A single user may however benefit in throughput if he or she doesn't respect the backoff. We call this player a malicious one.
- What happens if both players are malicious? If they are both correct?
- Are players better off playing one or other way?
- How would their strategy change if they now that the other is rational?
- And if they play the game several times?

Penalties

Left or Right?

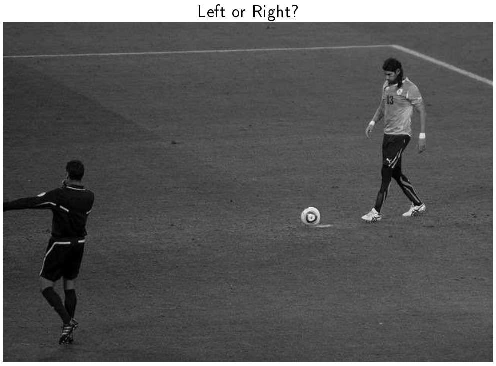

Notes:

- Consider a two-player situation, a goal keeper, and a penalty kicker
- Should the kicker kick right? left? middle?
- Should the goal keeper try left, right middle?
- How does this changes depending on the information of the players? And on their abilities?
- Real case analysis: [4]

Auctions

Electricity traded for delivery the following day in 24 hour intervals.

Notes:

- Data source:
  http://www.epexspot.com/en/market-data/dayaheadauction
  06/10/2017
- Orders contain up to 256 price/quantity combinations for each hour of the following day. Prices must be between - 500 euros/MWh and 3000 euros/MWh
- How much to bet?
- Which rules for the auction are more convenient for the bidders? for the auctioneer?
- How to set rules?

Building Powerlines

Who should pay the cost?

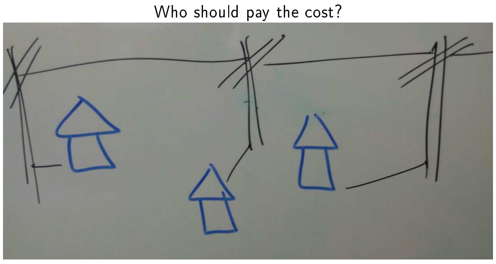

Notes:

- Consider a joint project, where players have benefits for acting together
- In this case, we consider the connection to the power grid of isolated houses, the connection solely of one of them must incur high costs, while the connection of all of them can profit from the connection of the other houses
- If they were to split the cost, how would they do it? Should each house pay its "last mile"? Should they equally shared the cost?
- Would they prefer to have their own connection or enter into the joint network?

And much more...

- **Transportation** The delay on a route depends on my choice of the route and the choice of every other driver
- **Internet, QoS, Wireless access** The throughput obtained depends on my choice and other user's choices
- **Smart Grids**
  - Recharge scheduling, the cost of the energy depends on the total grid load
  - Demand prediction e.g. for grid dimensioning
  - Energy trading (auctions)

Notes:

- Game theory has historically been a domain of economists
- Many economic applications such as sell or share of public goods, competition, price fixation, negotiation, etc.
- Also studied by sociologists and psychologist, since many behaviors found in the nature have been shown to fit the usually hypothesis of game theory
- Computer scientists have also, but more recently, study game theory, for distributed optimization, for instance
- Roughly we can say that economists care about the outcomes of situations, and computer scientists on how to get to them, how to compute them, how to find a good tradeoff between complexity and optimality of the outcome

Basic assumptions

- Rational players (pursue an objective)
- Strategically reasoning (use learning or assumptions about others' behaviors)

Notes:

- In order to predict the outcome of a situation, there are a set of common assumptions that are made
- this assumptions are more or less true depending on the context
- There is a rich literature about real experiments and comparisons with the outcomes predicted by the game theoretic approach
- For some cases theory is amazingly accurate with respect to reality
- For other cases less
- The basic assumptions are thus that players are rational and strategically reasoning

The Theory of Rational Choice

- Two elements: set of available actions, and preferences over the actions
- Rationality: a rational player will chose an action such that it is at least as good as any other available action according to its preferences.

Preferences

· The only assumption: if a player prefers a to b and b to c, then she also prefers a to c

· Can be also expressed by a pay-off function

Notes:

- Used by many game theoretical models
- Two elements: every player has a set of available actions, and preferences over those actions
- Preferences usually expressed by a payoff function, a function u that associates to every action a numerical value such that $u(a) > u(b)$ $\iff$ the player *prefers* a to b.
- Remark: for any strictly increasing function f, $f \circ u$ expresses the same preferences

Basic classification

- Non-cooperative
  - games in strategic form vs in extensive form
  - static games vs dynamic games, repeated games
  - games with perfect information vs with imperfect information
  - games with complete vs with incomplete information
- Cooperative
  - transferable utility, non-transferable utility, bargaining games

Notes:

- Several types of games, some of them are listed here
- These are some of the existing classifications
- There exist different solution concepts for different type of games
- We will talk of this term solution concept, what is it? It is just a solution that predicts the outcome of the game.

El curso

·Organización

·Fechas

·Evaluación

·Temario

Notes:

- Esta semana de 17h30 a 20h30
- Algunos ejercicios rápidos sobre el fin de la clase del tipo "quizz"
- Evaluación: lista de ejercicios + análisis de paper
- Objetivos del estudio de artículo: entender un paper científico publicado en una conferencia internacional o en una revista, realizar un resumen discutiendo fortalezas, debilidades del modelo y resultados. Realizar 4 slides resumiendo estas cosas. Las slides seran compartidas entre todos los asistentes. Trabajo individual.
- Fechas: entrega de ejercicios y análisis paper: fin de noviembre

Agenda

Notes:

Agenda

Notes:

Strategic Games (or Games in Normal Form)

Illustrative and classical example: the *Prisoner's Dilemma*

A \ B	confess	don't confess
confess	(3,3)	(0,4)
don't confess	(4,0)	(1,1)

Table: The situation: years of prison for each possible combination of actions

Notes:

- rational choice → choose utility functions that represent the preferences ordering
- this example represents situations where there is a gain from cooperation, but players have incentives to "free ride"
- many situations can be modeled with it, potentially with different payoffs but same preferences (e.g. TCP congestion, wireless access examples)

A \ B	confess	don't confess
confess	(1,1)	(4,0)
don't confess	(0,4)	(3,3)

Table: The game: Outcomes and Preferences for the different action profiles.

- What would be the outcome under the assumption of rational choice?

Definition

Formalization

Mathematically defined by the triplet (N, A, u)

. One-shot simultaneous, independent choices

Notes:

- Strategic games are non-cooperative games where players know all the game setting, and play simultaneously
- Mathematically specified by the triplet $G = (N, A, u)$ (set of players, set of actions for each player, preferences over *action profiles*)
- Def. action profile: the list of all the players' actions
- N set of $|N|$ players
- $A = \prod_{i \in N} A_i$, with A_i set of *actions* for each player
- u preferences expressed through *utilities* or pay-offs $u_i : A = A_1 \times \ldots \times A_i \times \ldots \times A_{|N|} \to \mathbb{R}$
- $u = \{u_i\}_{i \in N}$

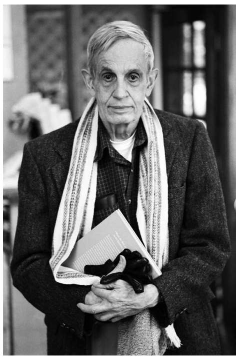

Notes:

- Photo: John Nash, By Peter Badge / Typos1 - OTRS submission by way of Jimmy Wales, CC BY-SA 3.0, `https://commons.wikimedia.org/w/index.php?curid=6977799`
- Very brief Nash's bio:
  - Nash's PhD thesis (Princeton University, 28 pages, he was 22 years old) introduces Nash equilibrium notion and result about existence of NE we shall see
  - At Princeton: Bargaining theory
  - Academic position in MIT
  - 1994 Nobel Prize in Economics Sciences joint with Harsany and Selten

Solution concept: Nash Equilibrium

- One of the most important solution concepts: Nash equilibrium (NE)

Nash Equilibrium

An action profile at which no player has any reason to unilaterally change his/her action choice.

Notes:

- By "no reason" we mean, no player will be better-off by unilaterally changing his/her action
- Mathematically, a NE is an action profile a^* such that $u(a^*) \geq u(a_i, a^*_{-i})$ for every a_i and for all i
- The actions as we have defined are also called *pure strategies*
- This definition of NE is also known as NE in pure strategies
- Strict Nash equilibrium implies no player will be better off or equal by unilaterally changing his/her action
- For strict NE, change $\geq$ for $>$

Example: the Prisoner's Dilemma

- Players: A, B
- Actions: confess, don't confess

		B: confess	B: don't confess
A	confess	(3,3)	(0,4)
A	don't confess	(4,0)	(1,1)

Table: Years of prison for each action profile.

		B: confess	B: don't confess
A	confess	(1,1)	(4,0)
A	don't confess	(0,4)	(3,3)

Table: Outcomes and Preferences.

Notes:

- Observe that the NE in this case (confess,confess) is not the optimal solution

Is the NE always unique?

The-battle-of-the-sexes game

		He: B	He: F
She	B	(1,2)	(0,0)
She	F	(0,0)	(2,1)

Notes:

- A NE might not be unique, example:

Does a NE always exist?

The Matching Pennies game

A \ B	T	H
T	(1,-1)	(-1,1)
H	(-1,1)	(1,-1)

Notes:

- A NE in pure strategies doesn't always exist, example:

Some facts about NE in pure strategies

- It does not always coincides with the optimal solution (e.g. The Prisoner's Dilemma)
- Not all games have a NE
- The NE can be not unique (several NE for the same game)

Notes:

- It does not always coincides with the optimal solution (e.g. The Prisoner's Dilemma)
- Not all games have a NE (e.g. Matching pennies)
- The NE can be not unique (several NE for the same game, e.g. BoS)

Searching for NE - Some definitions

Best response (BR)
For player i is a mapping that associates a set of actions for player i, to a list of other players' actions (a_{-i}), such that any action in that set is best response to a_{-i}.

NE, definition using BR functions
Is the action profile a^* is a NE iff every player's action is a BR to the other players actions

Notes:

- in complicated games, it is not easy to find the nash equilibria by simple inspection, as we did in the previous examples
- working with the best response function is sometimes easier
- the best response function is a set-valued function
- associates to any list of other players actions, a set of actions
- $BR_i(a_{-i}) = \{a_i : u(a_i, a_{-i}) \geq u(a_i', a_{-i}), \forall a_i' \in A_i\}$
- Nash equilibrium a^* iff $a_i^* \in BR_i(a_{-i}^*)$ for every player i.
- Method: find BR for each player, find the action(s) profile(s) that satisfy the NE definition
- Note that this is the intuitively approach one uses when searching on the table representation of the game
- Concerning complexity, the problem of finding the NE is not NP -hard (unless NP = coNP), it is something called "PPAD-hard"

Using BR to find NE, case of continuous set of actions

- Derivation can be used to determine the NE
- For two players 1 and 2: draw the best-response in terms $BR_1(a_2) = \text{argmax}_{a_1} u_1(a_1, a_2)$ and $BR_2(a_2) = \text{argmax}_{a_2} u_2(a_1, a_2)$.
- A Nash equilibrium is an intersection point of the best-response curves:

Notes:

- Example:

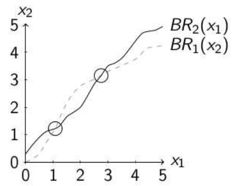

Dominant actions - definitions

Dominant actions

An action a_i **dominates** another action a'_i if i is "better off" with a_i than with a'_i, no matter what the other players do. An action is **dominant** if it dominates **all** actions.

- **Dominated actions:** In the previous case, we say that a'_i is **dominated** by a_i

Remarks:

- Can a dominated action be part of a NE?
- Weak domination, change "better" for "no worse" (with "better" for at least one set of other players' actions)

Notes:

- a_i **strictly dominates** a'_i if $u(a_i, a_-i) > u(a'_i, a_-i)$ for every a_{-i}
- a_i **weakly dominates** a'_i if $u(a_i, a_-i) \geq u(a'_i, a_-i)$ for every a_{-i} and a_i **strictly dominates** a'_i if $u(a_i, a_-i) > u(a'_i, a_-i)$ for some a_{-i}
- a strictly dominated action is never part of a NE, since a strictly dominated action is never best response to any list of other players actions
- A profile consisting on all dominant strategies is a NE
- A NE in strictly dominant strategies is unique
- In order to find NE we can proceed by eliminating strictly dominated actions

Example - Elimination of strictly dominated strategies

Player 2

Player 1		x	y	z
	a	(1,2)	(1,2)	(0,3)
	b	(4,0)	(1,3)	(0,2)
	c	(3,1)	(2,1)	(1,2)
	d	(0,2)	(0,1)	(2,4)

Table: IDSDS

Notes:

- Find all strategies that survive the iterative strictly dominated strategy elimination
- Verify that you found the NE
- Exercise originally proposed in Games, Strategies and Decision Making. Joseph Harrington Jr. Worth Publishers. (Second edition)

Pareto-optimum: another solution concept

Definition (Pareto-optimum)

An action profile is Pareto-optimal if there is no other action profile that can increase one player's utility without decreasing another player's utility.

- Is a NE necessarily a Pareto-Optimum?

Notes:

- Formally we say that an action profile a **Pareto dominates** an action profile a' if for all $i \in N$, $u_i(a) \geq u_i(a')$, and there exists some $j \in N$ for which $u_j(a) > u_j(a')$.
- We say that an action profile a is **Pareto optimal** if there does not exist another action profile $a' \in A$ that Pareto dominates a.
- A NE is not necessary Pareto-optimum (e.g. the Prisoner's Dilemma)
- In other words, at a pareto optimum, there is no way of making everyone happier

The Price of Anarchy: An Efficiency Measure

- The Price of Anarchy (PoA) measures the loss of efficiency by acting in a decentralized way
- A *Social Welfare* function must be defined W, mapping the actions profiles into a real
  - E.g. the sum of all player's utilities $W(a) = \sum_{i \in N} u_i(a)$, $a \in A$

$$PoA = \frac{\text{Social Optimum}}{\text{Worst Social Welfare at a NE}},$$

- Remark: $PoA = 1 \Rightarrow$ there is no loss of efficiency
- Exercise: Which is the PoA for the Prisoner's Dilemma if we consider $W(a) = \sum_{i \in N} u_i(a)$?

Notes:

- Welfare function is a function $W : A \to \mathbb{R}$
- The price of anarchy is then defined as $PoA = \frac{max_{a \in A} W(a)}{min_{a \in A^{NE}} W(a)}$, where A^{NE} is the set of Nash equilibria
- Which is the PoA for the Prisoner's Dilemma if we consider $W(a) = \sum_{i \in N} u_i(a)$? $PoA = \frac{6}{2} = 3$

Mixed strategies - definitions

- A mixed strategy for a player is a probability distribution on the set of his/her pure actions
- Utility u_i is a function whose expected value represents player i's preferences over the set of proba. distributions the set of actions profiles

Notes:

- Note that we talked about *actions*, now we talk about *strategies*
- Expected utility, $E_\pi[u_i] = \sum_{a \in A} u_i(a) \prod_j \pi_j(a_j)$
- Several interpretations of mixed strategies, not very intuitive concept
- For instance, models a situation where the participants' choices are not deterministic but are regulated by probabilistic rules
- Other interpretations:
- large population of players, each player chooses a pure action, and the payoff depends on the fraction of agents choosing each action. This represents the distribution of pure strategies (does not fit the case of individual agents).
- same game being played several times independently.

Preferences and expected utility

- Common assumption: Preferences are represented by the expected value of a payoff function. These are the so-called vNM preferences

John von vonNeumann (1903-1957, Hungry-US) By LANL [Public domain or Public domain], via Wikimedia Commons

Oskar Morgenrstern(1902 – 1977) , n.d., from the Oskar Morgenstern Papers.

Notes:

- A player has a preference over deterministic outcomes, eg he prefers a to b to c
- What are their preferences over non deterministic outcomes? i.e. over probability distributions over the deterministic outcomes?
- We need to add something to the model
- Common assumption: payoff (or utility function)
- we assume preferences over proba. distributions that can be represented by the expected value of a function over deterministic outcomes
- This function is usually called payoff or utility
- John von Neumann, the most important figure in the early development of game theory
- with economist Oskar Morgenstern wrote Theory of games and economic behavior, the book that established game theory as a field (around 1944)

NE in Mixed strategies - Definitions

NE in mixed strategies

is a set of **distribution functions** such that no user can unilaterally improve his expected utility by changing alone his/her distribution.

Notes:

- Formally, NE in mixed strategies is defined as a set of distribution π^* such that $\forall\, i, \forall\, \pi_i\ E_\pi[u_i] \geq E_{(\pi_i, \pi^*_{-i})}[u_i]$
- we can extend the definitions of best responses, dominant strategies etc.
- A best response in mixed strategies for player i is the set of player i's best mixed strategies when the other players' mixed strategies are given by π_{-i}

NE in Mixed strategies - Results

Theorem (Nash 1950)

Any finite n-person non-cooperative game has at least one equilibrium in mixed strategies.

Theorem (for finite games)

At equilibrium, for each player, any strategy yields the same expected utility.

Notes:

- Probed by Nash in his PhD thesis

Example: The Matching Pennies Game

At equilibrium,

- A plays H with probability x (and T with prob. $1 - x$)
- B plays H with probability y (and T with prob. $1 - y$)

		B: T	B: H
A	T	(1,-1)	(-1,1)
A	H	(-1,1)	(1,-1)

Find the NE. It might be helpful to remember that:

- At equilibrium, any action yields the same expected utility
- At equilibrium, all distributions are best responses

Notes:

- NE in mixed strategies:
- Hence, $E\{u_A(H)\} = E\{u_A(T)\}$ with $E\{u_A(H)\} = x.y.1 + x.(1-y).(-1)$ and $E\{u_A(T)\} = (1-x).y.(-1) + (1-x).(1-y).(1)$ and likewise for player B.
- Thus, the mixed strategy equilibrium is $((1/2, 1/2), (1/2, 1/2))$

Exercise: The Jamming game

		Jammer: channel 1	Jammer: channel 2
Regular Transmitter	channel 1	(-a,b)	(c,-d)
Regular Transmitter	channel 2	(c,-d)	(-a,b)

- a,b,c,d are all strictly positive

Notes:

- Two stations want to transmit in wireless, shared medium
- One is the regular transmitter, the other one is a jammer
- Payoffs are thus as in the table, where all values are strictly positive
- Observe that there is no NE in pure strategies
- Find a NE in mixed strategies, using best responses

A real world example

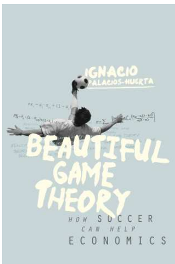

Notes:

- Ignacio Palacios-Huerta professor of management, economics, and strategy at the London School of Economics
- Paper: Professionals Play Minimax, Ignacio Palacios-Huerta, Review of Economic Studies (2003) 70, 395–415, 2003
- Model: strictly competitive games (2-players games with strategies diagonally opposed)
- Maxmin strategies: A maxmin mixed strategy for player i in a strategic game is a mixed strategy π_i^* that solves the problem $\max_{\pi_i} \min_{\pi_{-i}} E_{\pi_i^*, \pi_{-i}} u_i$
- In that case, If (π_1^*, π_2^*) is a Nash equilibrium then π_1^* is a maxminimizer for player 1, π_2^* is a maxminimizer for player 2, and $max_{\pi_1} min_{\pi_2} E_{(\pi_1, \pi_2)}[u_1] = min_{\pi_2} max_{\pi_1} E_{(\pi_1, \pi_2)}[u_1] = E_{(\pi_1^*, \pi_2^*)}[u_1]$.
- Intuition: analogous to the matching pennies, if you jump or kick randomly, with equal probability then the opponent cant exploit this strategy

Quizz on Lecture 1

Exercise 1. Quizz on game theory basics

A few questions for you to check that you have understood the notions of the course.

The questions proposed here are taken from the MOOC "Game Theory" on Coursera, created by Matthew Jackson, Kevin Leyton-Brown and Yoav Shoham. Many other interesting resources are available in that MOOC, including full lecture videos

Question 1 Consider the following normal-form game:

Player 1 Player 2	Movie	Theater
Movie	a, b	$0, 0$
Theater $0, 0$	c, d	

- N=1, 2
- A_i={Movie, Theater}. Each player chooses an action of either going to a movie or going to the theater.
- Player 1 prefers to see a movie with Player 2 over going to the theater with Player 2.
- Player 2 prefers to go to the theater with Player 1 over seeing a movie with Player 1.
- Players get a payoff of 0 if they end up at a different place than the other player.

Which restrictions should a, b, c and d satisfy?

a). $a > c, b > d$

b). $a > d, b < c$

c). $a > c, b < d$

d). $a < c, b < d$

Question 2 n people guess an integer between 1 and 100, and the winner is the player whose guess is closest to the mean of the guesses + 1 (ties broken randomly). Which of the following is an equilibrium?

a). All announce 1.

b). All announce 50.

c). All announce 75.

d). All announce 100.

Question 3 Consider the collective-action game:

Player 1 Player 2	Revolt	Not
Revolt	2,2	-1,1
Not	1,-1	0,0

When Player 1 plays "Not", for Player 2

a). "Revolt" is a best response.

b). "Not" is a best response.

c). "Revolt" and "Not" are both best responses.

d). There is no best response.

Question 4 Consider the following game in which two firms must decide whether to open a new plant or not:

Firm 1 Firm 2	Build	Not
Build	1,1	3,0
Not	0,3	2,2

Find all pure strategy Nash equilibria:

a). (Build, Not)

b). (Not, Not)

c). (Build, Build)

d). (Not, Build)

Question 5 Consider the game:

Player 1 Player 2	Left	Right
Up	2,1	1,1
Down	0,1	0,2

Which of the players has a strictly dominant strategy?

a). Player 1

b). Player 2

c). Both players

d). Neither player

Question 6 Consider the game:

Player 1 Player 2	Left	Right
Left	3,3	1,1
Right	1,4	1,1

Which of the following outcomes is Pareto-optimal? (There might be more than one, or none.)

a). (3,3)

b). (1,1)

c). (1,4)