Game Theory and Algorithms
Lake Como School of Advanced Studies,
7-11 September 2015 (Campione d'Italia)

Introduction to strategic games

Stefano Moretti

UMR 7243 CNRS
Laboratoire d'Analyse et Modélisation de Systèmes pour l'Aide à la décision (LAMSADE)
Université Paris-Dauphine

email: **stefano.moretti@dauhine.fr**

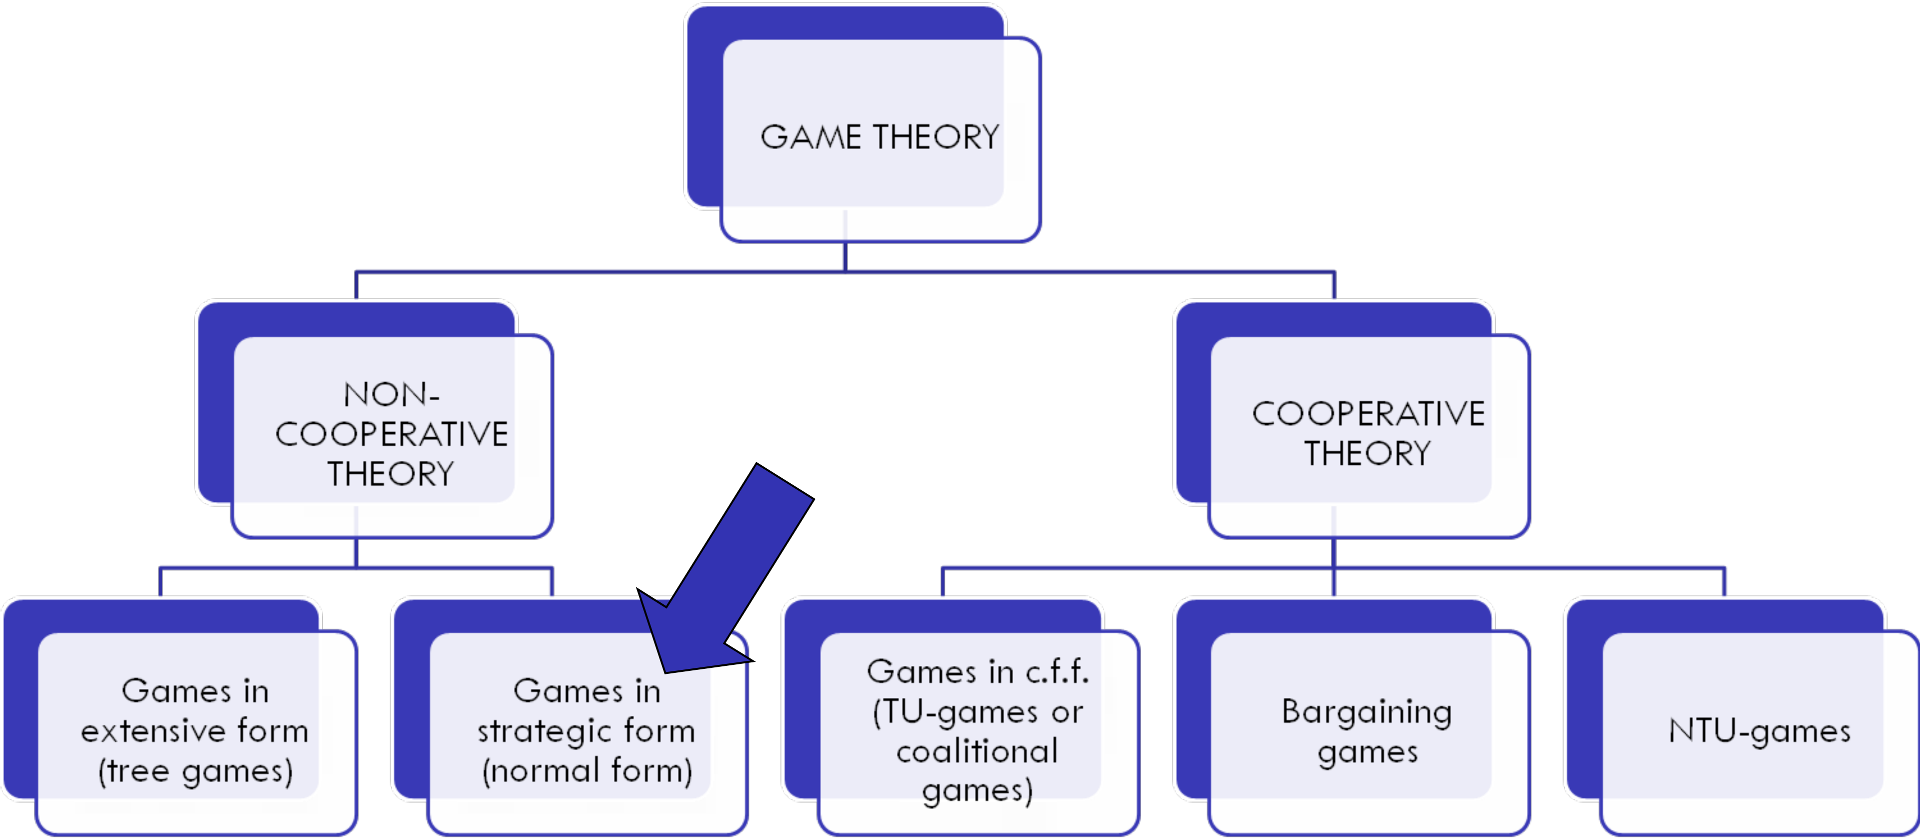

No binding agreements
No side payments
Goal: Optimal behaviour in conflict situations

binding agreements
side payments are possible (sometimes)
Goal: Reasonable sharing

Relevant characteristics

- Decision makers (=players) engaged in an interactive decision problem:
  - more than one decision maker (DM) (=player). [The “easy case", 1 DM, is left to Decision Theory (DT)]
  - the result is determined by the choices made by each player
  - the decision makers' preferences w.r.t. outcomes are (generally speaking) different.
- Classical assumptions about players: rational and intelligent.

Relevant parameters

players know the relevant data of the interaction decision problem:

- available strategies
- payoffs
- rationality and intelligence of each player
- each player knows that all players know what is listed above
- each player knows that each player knows...

COMMON KNOWLEDGE

- not available the possibility of binding agreements:

NON COOPERATIVE GAMES

From one to two DMs: Game form

- A game form (in strategic form), with two players, is:

$$(X, Y, E, h).$$

- New aspects w.r.t. decision theory:
- two DMs (we shall call them "players"), so two sets of available
- alternatives (choices, but here we use the word "strategies")
- $h : X \times Y \rightarrow E$ is the map that converts a couple of strategies

(one for each player) into an outcome.

- Easy to generalize to a finite set of players N:

$$(N, (X_i)_{i \in N}, E, h)$$

with $h : \prod_{i \in N} X_i \rightarrow E$.

Preferences of the players

- To get a game we need a second ingredient, the preferences of the players.
- If we have two players (called I and II), each will have his preferences: $\sqsupseteq_{I}$, $\sqsupseteq_{II}$.
- Each $\sqsupseteq_{I}$, $\sqsupseteq_{II}$ is a total preorder on E.
- We shall represent them by utility functions: u and v.
- We shall often make the assumption that these utilty functions are vNM (von Neumann-Morgenstern).

Game in strategic form

- Patching all together (game form + preferences)...
- We use utility functions. In the 2 players case:

$$(X, Y, E, h, u, v).$$

- The corresponding diagram:

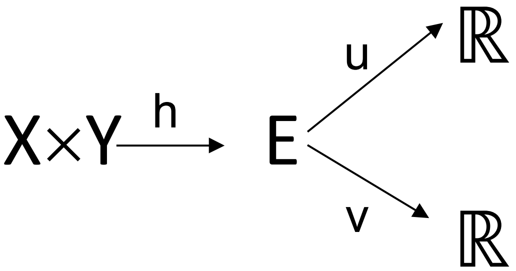

Game in strategic form: squeezed

- Still in the 2 players case:

$$(X, Y, f, g)$$

- where $f = u \circ h$ and $g = v \circ h$.
- The squeezed diagram:

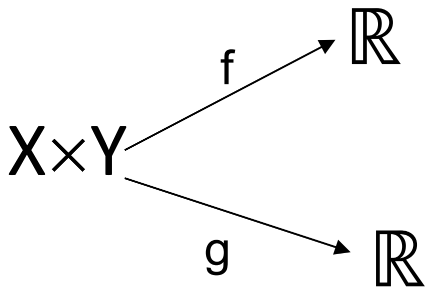

Example 1: Prisoner's dilemma

- Consider the following game:

I \ II	L	R
T	(3,3)	(1,4)
B	(4,1)	(2,2)

- You are the row player (I).
- The left number in each cell represents the evaluation that you give to the outcome. The number on the right represents the evaluation of player (II)…
- Which row do you choose? T or B?

Prisoner's dilemma tale (not very relevant)

- Two suspects are arrested by the police.
- The police have insufficient evidence for a conviction, and, having separated both prisoners, visit each of them to offer the same deal.
  - If one testifies (defects from the other) for the prosecution against the other and the other remains silent (cooperates with the other), the betrayer goes free and the silent accomplice receives the full 10-year sentence (strategies (B,L) or (T,R)).
  - If both remain silent, both prisoners are sentenced to only six months in jail for a minor charge (strategies (T,L)).
  - If each betrays the other, each receives a five-year sentence (strategies (B,R)).
- Each prisoner must choose to betray the other or to remain silent.
- Each one is assured that the other would not know about the betrayal before the end of the investigation.

Example 2: Coordination game

- Consider the following game:

I \ II	L	C	R
T	(0,0)	(1,1)	(0,0)
M	(0,0)	(0,0)	(1,1)
B	(1,1)	(0,0)	(0,0)

- Again you are the row player (I).
- Which row do you choose? T, M or B?

Looking for a solution

- What a player will/should do?
- “will”: the descriptive point of view. Aiming at predicting what players will do in the model and hence in the real game
- “should”: the normative point of view. Rationality is based on a teleological description of the players. Players have an “end” (not like apples, stones, molecules). So, they could do the “wrong” thing. We could give them suggestions.
- Can we say something on the basis of our assumptions?

Domination among strategies

- From decision theory we borrow the idea of domination among strategies:
- x_1 is (obviously) better than x_2 if:

$$h(x_1, y) \sqsupset h(x_{2,}\ y) \text{ for every } y \in Y$$

- We shall say that x_1 (strongly) dominates x_2.
- So, if x_1 dominates any other $x \in X$, then x_1 is the solution

Prisoner's dilemma

- The game is:

I \ II	L	R
T	(3,3)	(1,4)
B	(4,1)	(2,2)

- Obviously B and R are dominant strategies (for I and II respectively). So, we have the solution (B,R). Nice and easy.
- But... **the outcome is inefficient!**
- Both players prefer the outcome deriving from (T, L). And so? The problem is that players are (assumed to be) rational and intelligent.

Strategies to avoid

- A strategy which is (strongly) dominated by another one will not be played.
- So we can delete it. But then could appear new (strongly) dominated strategies for the other player. We can delete them.
- And so on...
- Maybe players are left with just one strategy each.
- Well, a new way to get a solution for the game.
- Technically: solution via iterated elimination of dominated strategies.

Strategies to avoid: example

I \ II	L	R
T	(2,1)	(1,0)
M	(1,1)	(2,0)
B	(0,0)	(0,1)

I \ II	L	R
T	(2,1)	(1,0)
M	(1,1)	(2,0)
~~B~~	~~(0,0)~~	~~(0,1)~~

I \ II	L	~~R~~
T	(2,1)	~~(1,0)~~
M	(1,1)	~~(2,0)~~

I \ II	L
T	(2,1)
~~M~~	~~(1,1)~~

Solution: (T, L)

Nash equilibrium

- Basic solution concept, for games in strategic form.
- (2 players only) Given G = (X, Y, f, g), $(x^*,y^*) \in X \times Y$ is a Nash equilibrium for G if:

$$f(x^*,y^*) \geq f(x,y^*) \text{ for all } x \in X$$

$$g(x^*,y^*) \geq g(x^*,y) \text{ for all } y \in Y$$

Interpretation: x^* is a best reply (max utility f) for player I when player II plays strategy y^*, and y^* is a best reply (max utility g) for player II when player I plays strategy x^*.

Nash equilibrium: examples (2)

NE calculation in BoS

Best reply for player I: max utility

Fix this strategy for II

I \ II	L	R
T	(2,1)	(0,0)
B	(0,0)	(1,2)

Best reply for player I: max utility

Fix this strategy for II

I \ II	L	R
T	(2,1)	(0,0)
B	(0,0)	(1,2)

Fix this strategy for I

Best reply for player II: max utility

I \ II	L	R
T	(2,1)	(0,0)
B	(0,0)	(1,2)

Fix this strategy for I

Best reply for player II: max utility

I \ II	L	R
T	(2,1)	(0,0)
B	(0,0)	(1,2)

Couples of strategies with both payoffs in red are N.E.

Nash equilibrium: examples (3)

- Example (battle of the sexes, BoS): (T, L) and (B, R) are Nash Equilibria (N.E.). Not unique!

I \ II	L	R
T	(2,1)	(0,0)
B	(0,0)	(1,2)

Nash equilibrium: examples (4)

Consider the following game (coordination game):

I \ II	L	C	R
T	(0,0)	(1,1)	(0,0)
M	(0,0)	(0,0)	(1,1)
B	(1,1)	(0,0)	(0,0)

(B, L), (T, C) and (M,R) are N.E.

Nash equilibrium: not unique

The battle of the sexes and the coordination game (and many others) have more than one NE.

- BIG ISSUE.
- players may have different (opposite) preferences on the equilibrium outcomes (see BoS)
- it is not possible to speak of equilibrium strategies. In the BoS, T is an equilibrium strategy? Or B?

One more problem

Example: matching pennies (MP)

I \ II	L	R
T	(-1,1)	(1,-1)
B	(1,-1)	(-1,1)

- There is no equilibrium?
- But Nash is famous (also) because of his existence thm (1950).
- But MP is a zero-sum game. So, even vN (1928) guarantees that it has an equilibrium.
- Where do we find it? Usual math trick: extend ($\mathbb{N}$ to $\mathbb{Z}$, sum to integral, solution to weak solution).

Mixed strategies

- The basic idea is that the player does not choose a strategy, but a probability distribution on strategies.
- Example: I have an indivisible object and I must assign it in a fair way to one of my children. It is quite possible that the best solution is to decide to assign it randomly (with a uniform probability distribution).

Mixed extension of a game

- Let's apply it to games in strategic form.
- Given a game G = (X, Y, f, g), assume X, Y are finite, and let

$$X = \{x_1, \ldots, x_m\},\ Y = \{y_1, \ldots, y_n\}.$$

- The mixed extension of G is $\Gamma = (\Delta(X), \Delta(Y), f', g')$, where:
- $\Delta(X)$ ($\Delta(Y)$)is the set of probability distributions on X (Y). An element of $\Delta(X)$ is

$p = (p_1, \ldots, p_m) \in \mathbb{R}^{m,}$ where p_i is the probability to play strategy x_i and…

$$f'(p, q) = \sum_{i \in \{1,\ldots m\}} \sum_{j \in \{1,\ldots n\}} p_i q_j f(x_i, y_j)$$

$$g'(p, q) = \sum_{i \in \{1,\ldots m\}} \sum_{j \in \{1,\ldots n\}} p_i q_j g(x_i, y_j)$$

- Of course, $(p,q) \in \Delta(X) \times \Delta(Y)$
- Notice that Γ is itself a game in strategic form. So, no need to redefine concepts (in particular, N.E.).

Interpretation?

- Of course, there is no mathematical problem in the definition of Γ.
- But: f’ and g’ can still be interpreted as payoffs for the players?
- The answer is YES if the original f and g are vNM utility functions. Otherwise, we cannot attach a meaning to the operations that brought us from G to Γ.

Mixed extension and equilibria for BoS

The BoS is:

		q	1-q
	I \ II	L	R
p	T	(2,1)	(0,0)
1-p	B	(0,0)	(1,2)

- Instead of using $((p_1,p_2),(q_1,q_2))$ we use $((p,1-p),(q,1-q))$, with $p,q \in [0,1]$. So:

  $$f'((p,1-p),(q,1-q))=2pq+1(1-p)(1-q)=(3q-1)p+(1-q)$$

- Given q, the best reply for player I to q is p* such that
  - $p^*=0$ if $0 \le q < 1/3$
  - $p^* \in [0,1]$ if $q=1/3$
  - $p^*=1$ if $1/3 < q \le 1$

Mixed extension and equilibria for BoS (2)

Given p, the best reply for player II to p* is

$q^*=0$ if $0 \leq p < 2/3$

$q^* \in [0,1]$ if $p=2/3$

$q^*=1$ if $2/3 < p \leq 1$

From the following picture we see there are 3 N.E.

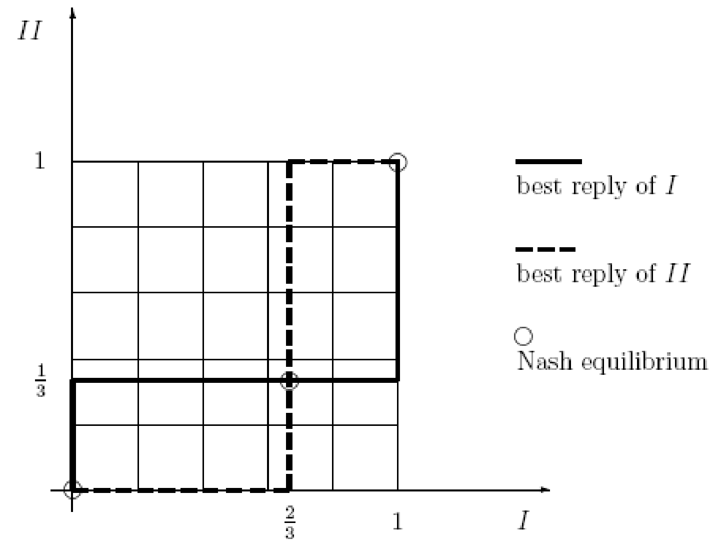

Potential games

A strategic game G = (X, Y, f, g) is said to be an (exact) potential game if

there exists $P: X \times Y \to \mathbb{R}$ such that:

for each $x_1, x_2 \in X$ and each $y \in Y$ $P(x_1,y)-P(x_2,y)=f(x_1,y)-f(x_2,y)$

for each $x \in X$ and each $y_1, y_2 \in Y$ $P(x,y_1)-P(x,y_2)=g(x,y_1)-g(x,y_2)$

P is said to be **a potential** for G.

Nash equilibria in potential games

Given a potential game G = (X, Y, f, g) it is obvious from the definition that (x*,y*) is NE for G iff (x*,y*) is a NE for (X,Y,P,P).

Theorem (obvious)
If (x*,y*) maximizes P then (x*,y*) is a NE
Corollary
A finite game with potential has a NE in pure strategy

Computationally interesting: it reduces the search for a NE to a search for maximum (or for minimum, if we have costs in the matrix instead of gains).

Nash equilibria in potential games

Remark: there are NE that are not potential maximizers, so we are not sure to find all NE.

I \ II	L	R
T	(1,1)	(0,0)
B	(0,0)	(0,0)

Potential maximizers can be seen as refinement of NE.

Congestion games

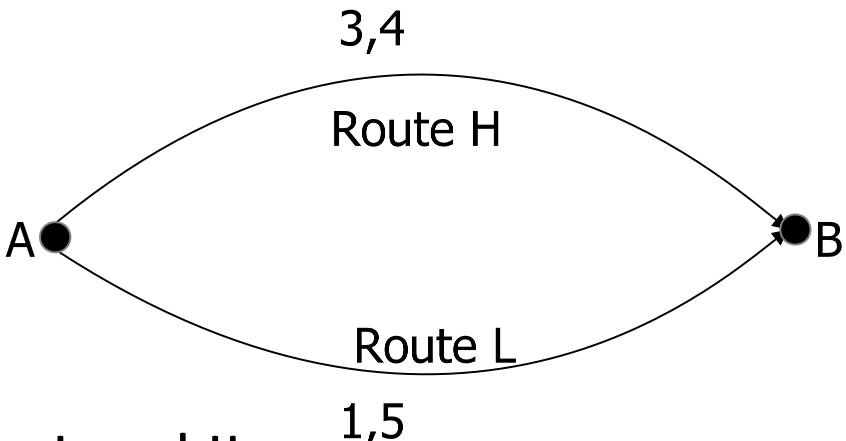

Two players I and II:

3 is the cost for each player to use the road H when only one player uses that road;

4 is the cost for each player to use the road H when precisely two players use that road.

Similar for route L.

Congestion games

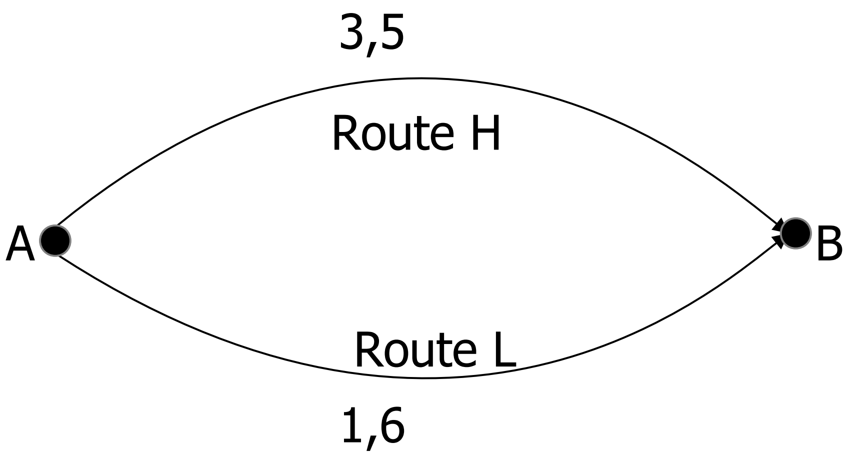

Values are costs!

I \ II	H	L
H	(5,5)	(3,1)
L	(1,3)	(6,6)

➡

potential

I \ II	H	L
H	8	4
L	4	7

All potential games are congestion games and vice-versa (Rosenthal (1973))

Ordinal potential games

From the definition of exact potential game (for player I)

for all $x_1, x_2 \in X$, $y \in Y$, $P(x_1,y)-P(x_2,y)=f(x_1,y)-f(x_2,y)$ (1)

From (1) it follows:

for all $x_1, x_2 \in X$, $y \in Y$, $P(x_1,y)>P(x_2,y)$ iff $f(x_1,y)>f(x_2,y)$ (2)

We cannot go back from (2) to (1). It is a weaker condition. So we speak of ordinal potential game. Notice that we can rewrite (2) using only preferences.

for all $x_1, x_2 \in X$, $y \in Y$, $(x_1,y)R(x_2,y)$ iff $(x_1,y) \sqsupseteq_I (x_2,y)$

for all $x \in X$, $y_1, y_2 \in Y$ $(x,y_1)R(x,y_2)$ iff $(x,y_1) \sqsupseteq_{II} (x,y_2)$

Where R, $\sqsupseteq_I$ and $\sqsupseteq_{II}$ are preferences induced on $X \times Y$.

Better-response dynamic (BRD)

BRD is a straightforward procedure by which players search for a NE of a game. Specifically:

While the current outcome **s** is not a NE:

Take an arbitrary player *i* and an arbitrary beneficial deviation s'_i for player *i* and move to the outcome $(s'_i, \mathbf{s}_{-i})$

BRD can only halt at a Pure NE

BRD cycles in any game without a NE.

BRD can also cycle in games that have a NE (see next slide)

For potential games the BRD terminates.

Better-response dynamic

I \ II	L	C	R
T	(1,-1)	(-1,1)	(-1,-1)
M	(-1,1)	(1,-1)	(-1,-1)
B	(-1,-1)	(-1,-1)	(-1,-1)

(B,R) is the unique NE