Silvia Rossi

Non-cooperative game theory

Lezione n. 9

Corso di Laurea:
Informatica

Insegnamento:
Sistemi
multi-agente

Email:
silrossi@unina.it

A.A. 2014-2015

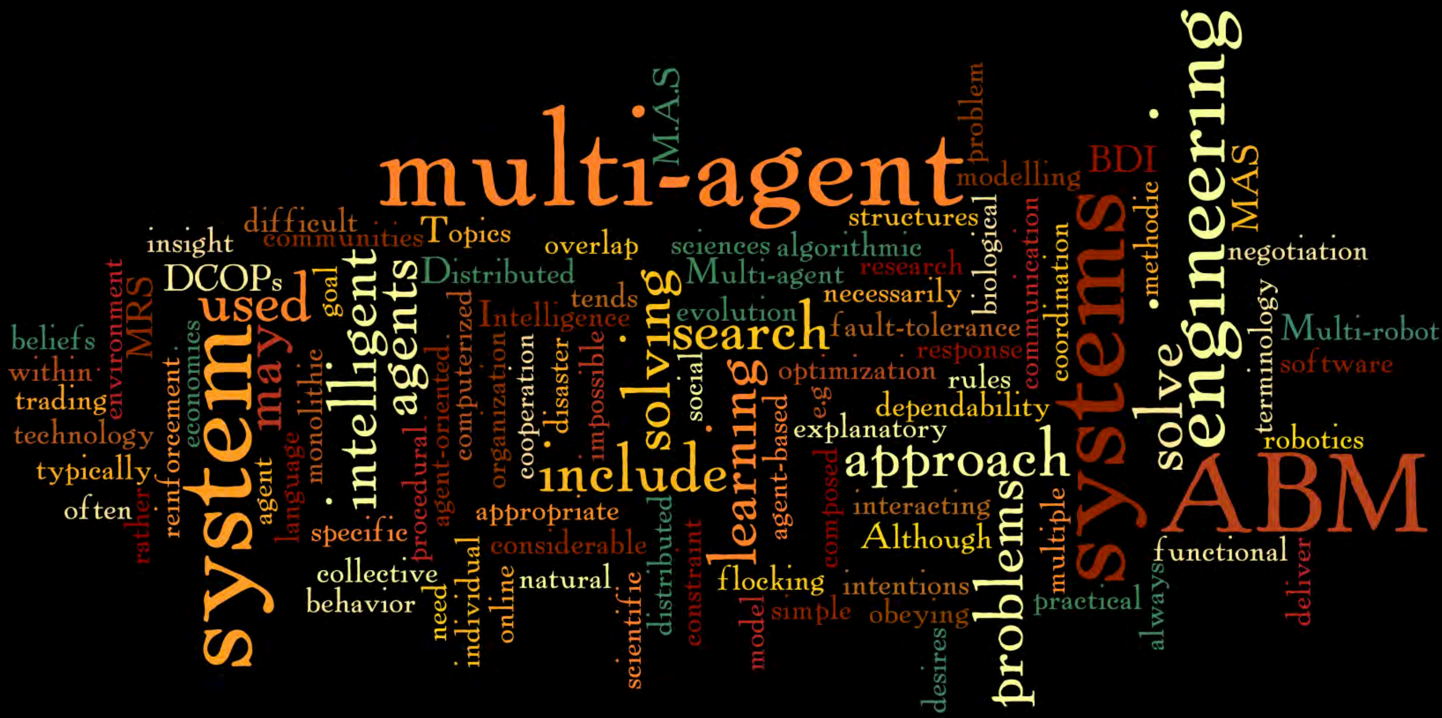

Distributed Rational Decision Making – Game Theory

(W: 3, MAS: 3.3)

Strategies in normal-form games

- Selection of one action -> **pure strategy**

- Players could randomize over the set of available actions according to some probability distribution -> **mixed strategy**

The **support** of a mixed strategy support of a s_i for a player i is the mixed strategy set of pure strategies $\{a_i | s_i(a_i) > 0\}$

Given a normal-form game (N,A, u), the **expected utility** u_i for player i of the mixed-strategy profile $s = (s_1, . . . , s_n)$ is defined as:

$$u_i(s) = \sum_{a \in A} u_i(a) \prod_{j=1}^{n} s_j(a_j).$$

How will a rational agent will behave in any given scenario? Play. . .

- Strategies that maximise social welfare;
- Pareto optimal strategies;
- Nash equilibrium strategy;
- Dominant strategy.

The social welfare of an outcome ω is the sum of the utilities that each agent gets from ω:

$$\sum_{i \in Ag} u_i(\omega)$$

Think of it as the "total amount of money in the system".

As a solution concept, may be appropriate when the whole system (all agents) has a single owner (then overall benefit of the system is important, not individuals).

Payoff Matrices

We can characterize the previous scenario in a *payoff matrix*:

Agent *i* is the *column player*

Agent *j* is the *row player*

		i	
		defect	coop
j	defect	1 1	4 1
	coop	1 4	4 4

An outcome is said to be *Pareto optimal* (or *Pareto efficient*) if there is no other outcome that makes one agent *better off* without making another agent *worse off*.

If an outcome is Pareto optimal, then at least one agent will be reluctant to move away from it (because this agent will be worse off).

Pareto domination

Strategy profile s Pareto Pareto dominates strategy profile s′ if for all $i \in N$, $u_i(s) \geq u_i(s')$, and there exists some $j \in N$ for which $u_j(s) > u_j(s')$

Pareto optimality

Strategy profile s is Pareto optimal if there does not exist another strategy profile $s' \in S$ that Pareto dominates s

If an outcome ω is *not* Pareto optimal, then there is another outcome ω′ that makes *everyone* as happy, if not happier, than ω.

“Reasonable” agents would agree to move to ω′ in this case.

(Even if I don’t directly benefit from ω′, you can benefit without me suffering.)

Every game must have at least one such optimum

Some games will have multiple optima

	C	D
C	-1,-1	-4,0
D	0,-4	-3,-3

	C	D
C	-1,-1	-4,0
D	0,-4	-3,-3

	C	D
C	1,1	0,0
D	0,0	1,1

	C	D
C	-1,-1	-4,0
D	0,-4	-3,-3

	C	D
C	1,1	0,0
D	0,0	1,1

	B	S
B	2, 1	0,0
S	0,0	1,2

	C	D
C	-1,-1	-4,0
D	0,-4	-3,-3

	C	D
C	1,1	0,0
D	0,0	1,1

	B	S
B	2, 1	0,0
S	0,0	1,2

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

- If you knew what everyone else was going to do, it would be easy to pick your own action

  - Let $a_{-i} = (a_1,\ldots,a_{i-1},a_{i+1},\ldots,a_n)$.
  - now $a = (a_{-i},a_i)$

Best response

a^*_i in $BR(a_{-i})$ iff $\forall a_i$ in A_i, $u_i(a^*_i,a_{-i}) \geq u_i(a_i,a_{-i})$

The best response is not necessarily unique.

When the support of a best response a^* includes two or more actions, the agent must be indifferent among them

- Now let's return to the setting where no agent knows anything about what the others will do
- What can we say about which actions will occur?

Idea: look for **stable action profiles.**

a = <a1, … , an> is a Nash equilibrium iff

$\forall i$, a_i in $BR(a_{-i})$.

$\forall i \in N \; \forall a_i \in A_i \; (a^*_{-1}, a^*_i) \geq_i (a^*_{-1}, a_i)$

So: no player i can improve in a*, if all the other players keep playing a^*_{i-1}

Dato $G = \langle N, (A_i), (\geq_i) \rangle$

$a^* \in a = a_1 \times a_2 \times \ldots \times a_n$ is Nash equilibrium if

$\forall i \in N\ \forall a_i \in A_i\ (a^*_{-1}, a^*_i) \geq_i (a^*_{-1}, a_i)$

So: no player i can improve in a*, if all the other players keep playing a^*_{i-1}

In general, we will say that two strategies s_1 *and* s_2 are in Nash equilibrium if:

1. under the assumption that agent i plays s_1, agent j can do no better than play s_2; and
2. under the assumption that agent j plays s_2, agent i can do no better than play s_1.

Neither agent has any incentive to deviate from a Nash equilibrium

Unfortunately:

1. Not every interaction scenario has a Nash equilibrium

2. Some interaction scenarios have more than one Nash equilibrium

Example: BoS

$N = \{1,2\}$

$A_1 = \{B,S\}$

$A_2 = \{B,S\}$

u_1, u_2

B: Bach

S: Strawinsky

Battle of the Sexes

	B	S
B	2,1	0,0
S	0,0	1,2

Play this game with someone near you. Then find a new partner and play again. Play five times in total.

Example: BoS (N.eq)

$N = \{1,2\}$

$A_1 = \{B,S\}$

$A_2 = \{B,S\}$

u_1, u_2

B: Bach

S: Strawinsky

Battle of the Sexes

	B	S
B	2,1	0,0
S	0,0	1,2

Example: BoS (N.eq)

$N = \{1,2\}$

$A_1 = \{B,S\}$

$A_2 = \{B,S\}$

u_1, u_2

B: Bach

S: Strawinsky

two equilibria:

(bach,bach) en

(strawinsky, strawinsky)

	B	S
B	2,1	0,0
S	0,0	1,2

Ex: coordination game

Mozart of Mahler?

Equal preferences

	Mo	Ma
Mo	2,2	0,0
Ma	0,0	1,1

Mozart of Mahler?

Equal preferences

	Mo	Ma
Mo	2,2	0,0
Ma	0,0	1,1

Ex: coordination game

Mozart of Mahler?

Equal preferences

two equilibria:

(Mozart,Mozart) and

(Mahler,Mahler)

N.eq right concept?

	Mo	Ma
Mo	2,2	0,0
Ma	0,0	1,1

consider cooperative game G

(2,2) seems the optimal solution

security strategy of 1 is r, gives 1!

Nash equilibria?

	L	R
l	2,2	0,0
r	1,1	1,1

consider cooperative game G

(2,2) seems the optimal solution

security strategy of 1 is r, gives 1!

Nash equilibria?

	L	R
l	2,2	0,0
r	1,1	1,1

m x n matrix

1 has strategies s1 and s2,

2 has t1, t2 and t3

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

m x n matrix

1 has strategies s1 and s2, 2 has t1, t2 and t3

Nash equilibrium (σ,τ):

$\forall s,t\ \pi_1(\sigma,\tau) \geq \pi_1(s,\tau)$

$\forall s,t\ \pi_2(\sigma,\tau) \geq \pi_2(\sigma,t)$

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

Given any particular strategy (either *C* or *D*) of agent i, there will be a number of possible outcomes

We say s_1 *dominates* s_2 if every outcome possible by i playing s_1 is preferred over every outcome possible by i playing s_2

A rational agent will never play a dominated strategy

So in deciding what to do, we can *delete dominated strategies*

Unfortunately, there isn't always a unique undominated strategy

strategy s_d of 1 strongly dominates s_i:

$\forall t\ \pi_1(s_d,t) > \pi_1(s_i,t)$

strategy s_d of 1 strongly dominates s_i:

$\forall t\ \pi_1(s_d,t) > \pi_1(s_i,t)$

s_2 strongly dominates s_1

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

strategy s_d of 1 strongly dominates s_i:

$\forall t\ \pi_1(s_d,t) > \pi_1(s_i,t)$

And weak if:

$\forall t\ \pi_1(s_d,t) \geq \pi_1(s_i,t)$

$\exists t\ \pi_1(s_d,t) > \pi_1(s_i,t)$

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

strategy s_d of 1 strongly dominates s_i:

$\forall t\ \pi_1(s_d,t) > \pi_1(s_i,t)$

And weak if:

- $\forall t\ \pi_1(s_d,t) \geq \pi_1(s_i,t)$
- $\exists t\ \pi_1(s_d,t) > \pi_1(s_i,t)$

t_1 weakly dominates t_2

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

s_2 strongly dominates s_1

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

s_2 strongly dominates s_1

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

s_2 strongly dominates s_1

t_1 weakly dominates t_2

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

s_2 strongly dominates s_1

t_1 weakly dominates t_2

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

s_2 strongly dominates s_1

t_1 weakly dominates t_2

In the new game t_1 and t_3 are not weakly dominated

(s_2,t_1) and (s_2,t_3) N.eq!

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

	t_1	t_2	t_3
s_1	1,1	2,0	3,-1
s_2	2,0	4,0	6,0

A strategy is strictly (weakly) dominant *for an agent if it strictly (weakly) dominates any other strategy for that agent.*

It is obvious that a strategy profile (s1, . . . , sn) in which every si is dominant for player i (whether strictly, weakly) is a Nash equilibrium.

An equilibrium in strictly dominant strategies is necessarily the unique Nash equilibrium.

	C	D
AE	2,0	1,1
AF	0,2	1,1
BE	3,3	3,3
BF	3,3	3,3

	C	D
AE	2,0	1,1
AF	0,2	1,1
BE	3,3	3,3
BF	3,3	3,3

	C	D
AE	2,0	1,1
AF	0,2	1,1
BE	3,3	3,3
BF	3,3	3,3

equilibrium gone!

	C	D
AE	2,0	1,1
AF	0,2	1,1
BE	3,3	3,3
BF	3,3	3,3

	C	D
AE	2,0	1,1
AF	0,2	1,1
BE	3,3	3,3
BF	3,3	3,3

	C	D
AE	2,0	1,1
AF	0,2	1,1
BE	3,3	3,3
BF	3,3	3,3

The elimination order does not matter when we remove *strictly dominated strategies (Church–Rosser property).*

With weakly dominated strategies:

- subgameperfect equilibrium can be lost
- Order of elimination matters

Order of elimination

	C	D
AE	2,0	1,1
AF	0,2	1,1
BE	3,3	3,3
BF	3,3	3,3

Two men are collectively charged with a crime and held in separate cells, with no way of meeting or communicating. They are told that:

- if one confesses and the other does not, the confessor will be freed, and the other will be jailed for three years
- if both confess, then each will be jailed for two years

Both prisoners know that if neither confesses, then they will each be jailed for one year

Payoff matrix for prisoner’s dilemma:

		i	
		defects	**coperates**
j	**defects**	**2** **2**	**0** **5**
	coperates	**5** **0**	**3** **3**

Top left: If both defect (confess), then both get punishment for mutual defection

Top right: If i cooperates and j defects, i gets payoff of 0, while j gets 5

Bottom left: If j cooperates and i defects, j gets payoff of 0, while i gets 5

Bottom right: Reward for mutual cooperation (not confessing)

The *individual rational* action is *defect*
This guarantees a payoff of no worse than 2, whereas cooperating guarantees a payoff of at most 0

So defection is the best response to all possible strategies: both agents defect, and get payoff = 2

But *intuition* says this is *not* the best outcome: Surely they should both cooperate and each get payoff of 3!

l'*equilibrio di Nash* rappresenta quindi la situazione nella quale il gruppo si viene a trovare se **ogni componente del gruppo fa ciò che è meglio per se.**

l'ottimo di Pareto è razionale dal punto di vista collettivo, ma non lo è affatto dal punto di vista individuale.

(*D,D)* is the only Nash equilibrium.

All outcomes *except (D,D)* are Pareto optimal.

(*C,C)* maximises social welfare.

l'equilibrio di Nash può non essere un ottimo di Pareto