A tutorial

Taking rational decisions: An Introduction to Game Theory

Georgios Chalkiadakis

Associate Professor

School of Electrical and Computer Engineering

Technical University of Crete

Game Theory

- a branch of economics that deals with decision-making in environments full of self-interested entities
- provides tools for the strategic considerations of intelligent, “rational” agents

Putting Things into Perspective

It's not a brain…it's a social network…part of the WWW… or is it our collective brain?

It's not a brain…it's a social network…part of the WWW… or is it our collective brain?

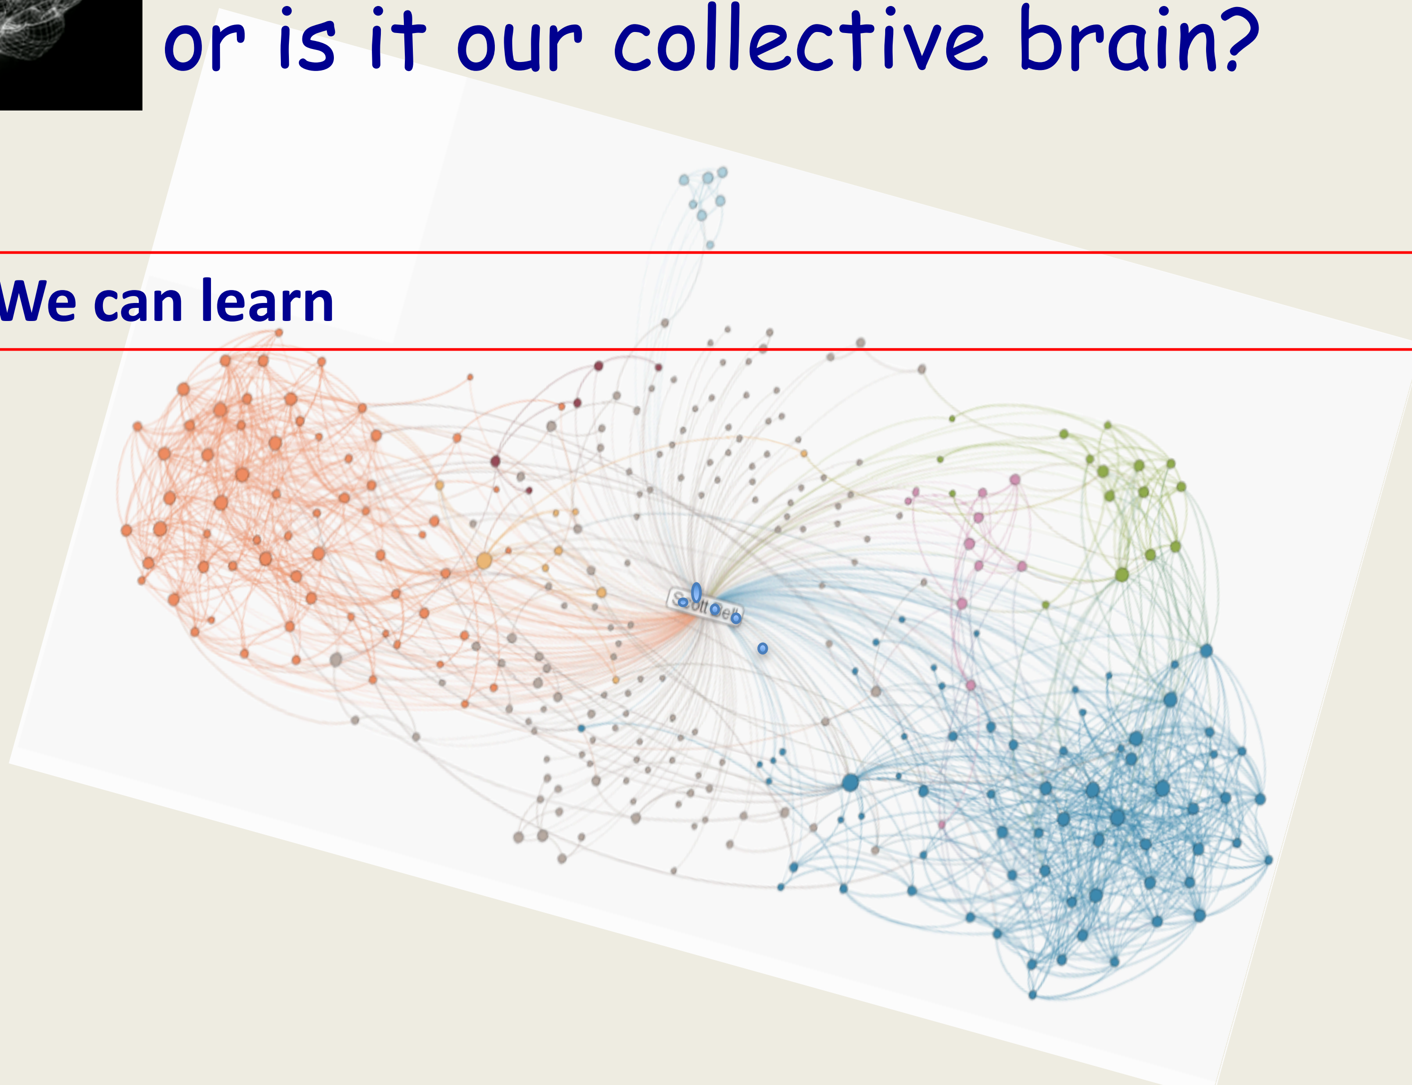

✓ **We can learn**

It's not a brain…it's a social network…part of the WWW… or is it our collective brain?

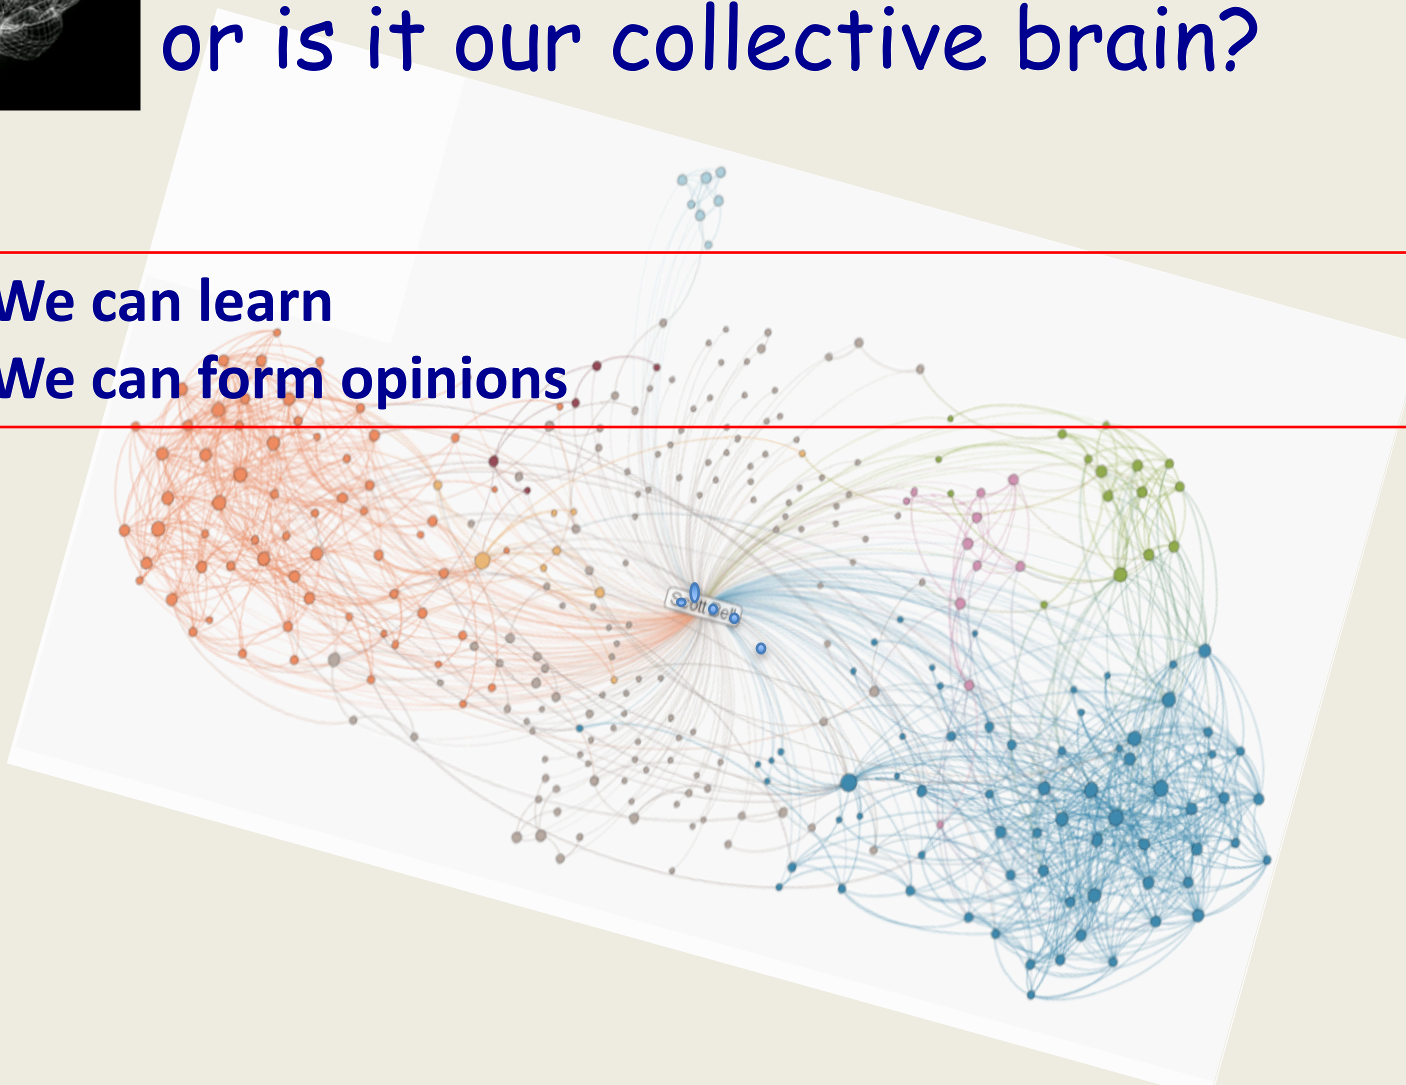

- ✓ **We can learn**
- ✓ **We can form opinions**

It's not a brain…it's a social network…part of the WWW… or is it our collective brain?

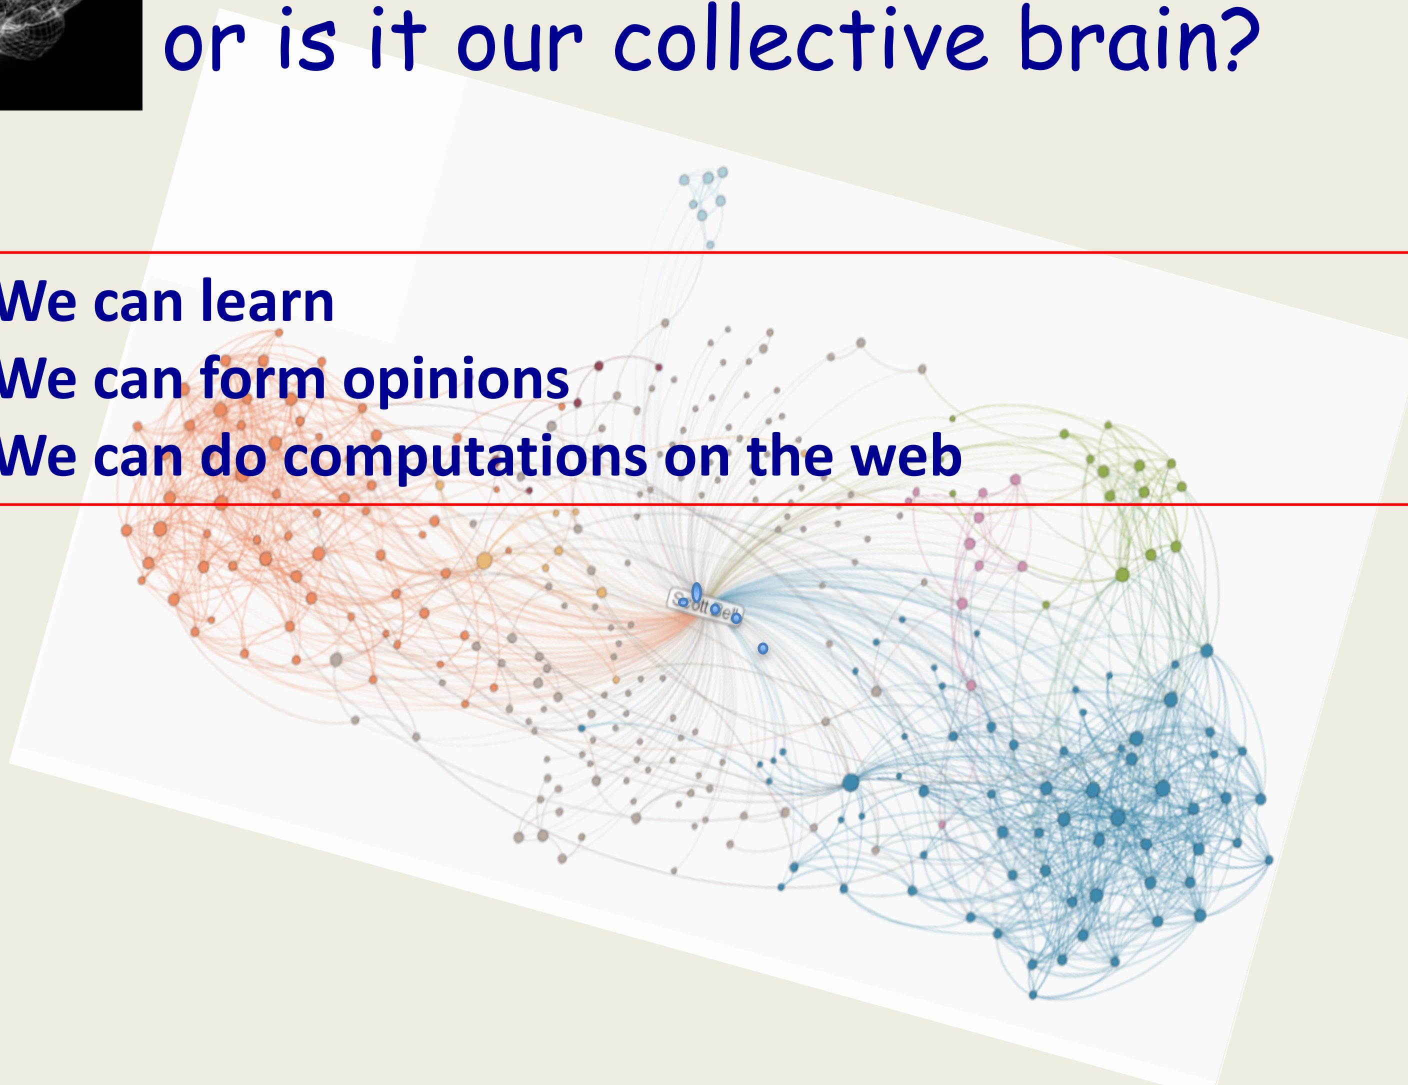

- ✓ **We can learn**
- ✓ **We can form opinions**
- ✓ **We can do computations on the web**

It's not a brain…it's a social network…part of the WWW… or is it our collective brain?

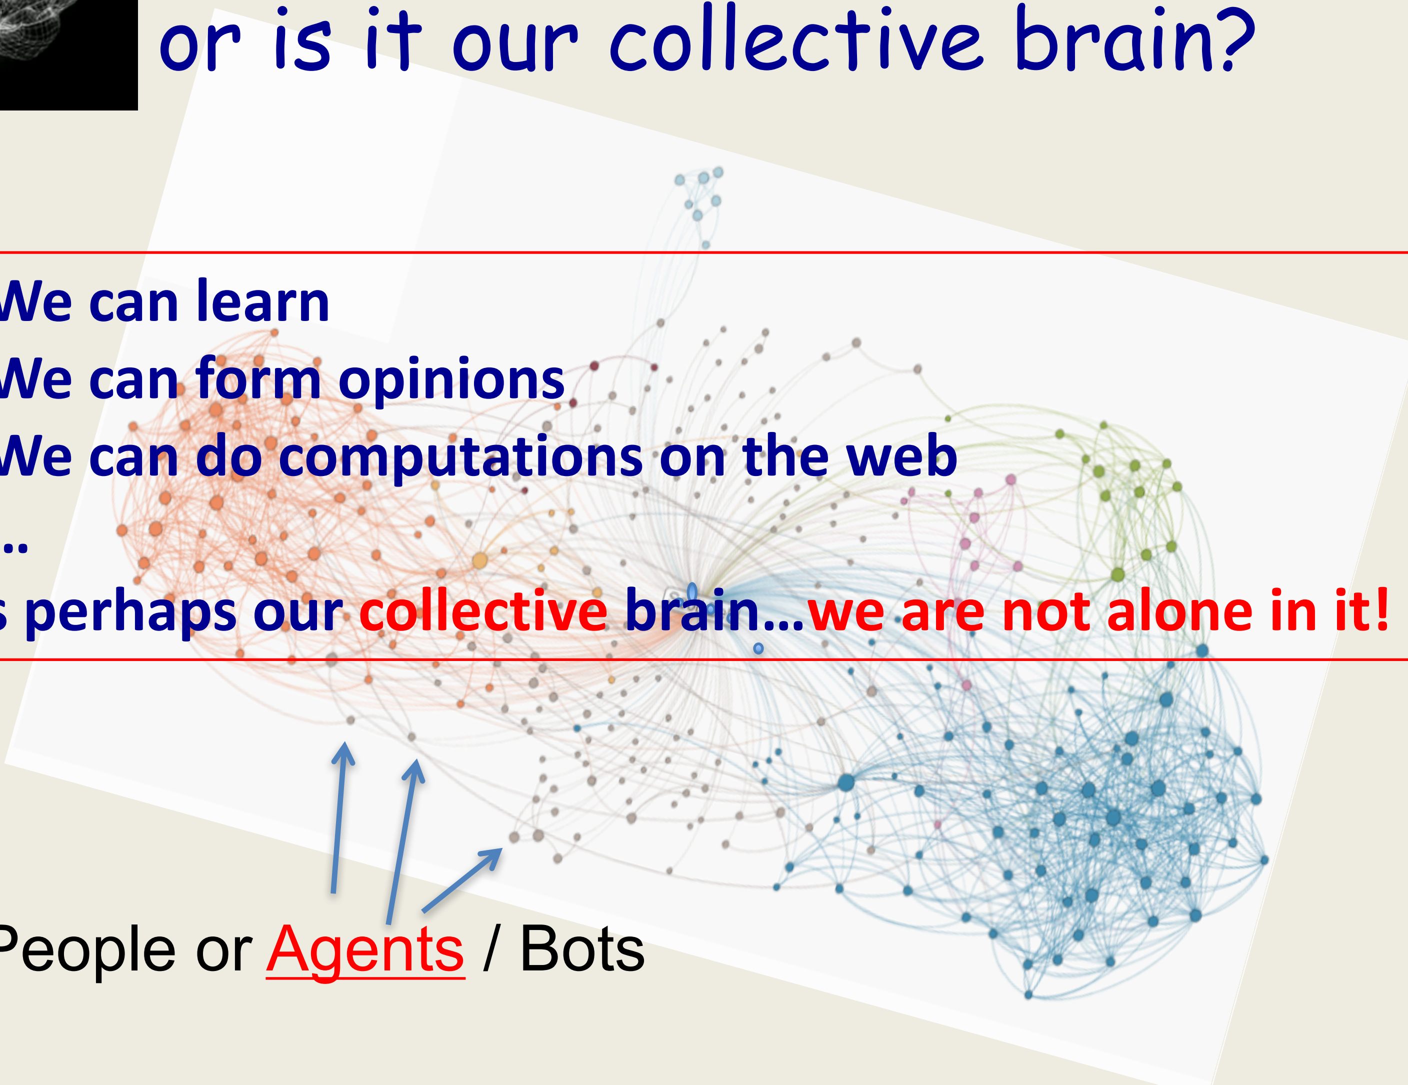

- ✓ **We can learn**
- ✓ **We can form opinions**
- ✓ **We can do computations on the web**
- ✓ **…**

It is perhaps our collective brain…we are not alone in it!

People or Agents / Bots

It's not a brain…it's a social network…part of the WWW… or is it our collective brain?

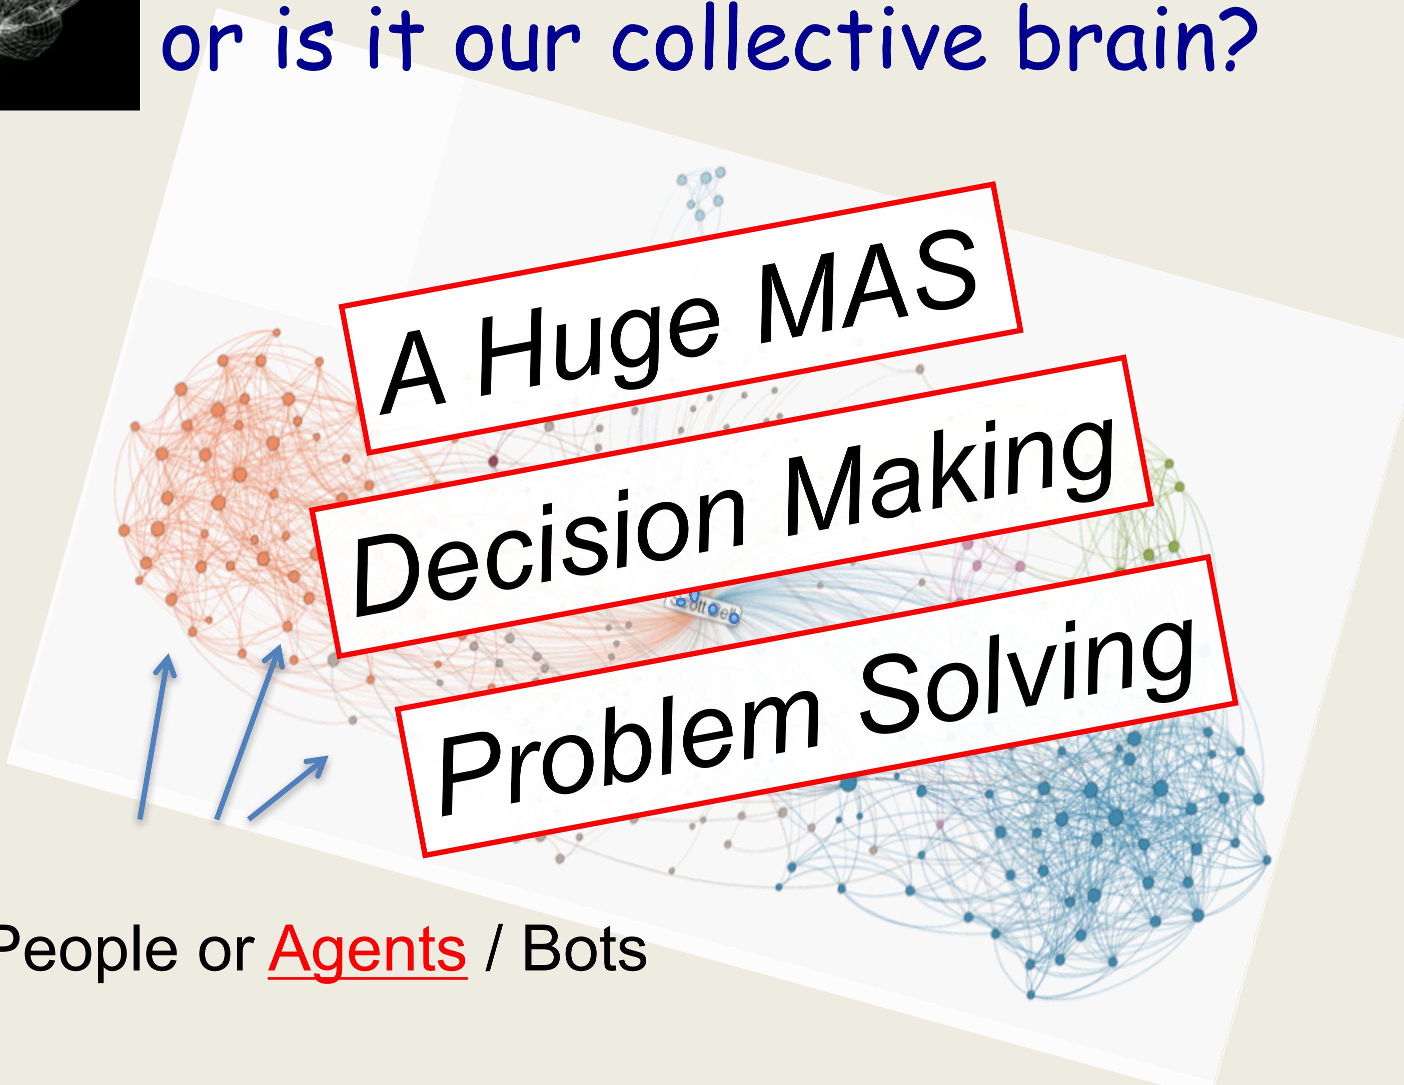

People or Agents / Bots

Who are these agents?

- Rational Agent
  - Reactive
  - Proactive
  - **Socially able**
  - Acts to achieve the best possible outcome, given facts/knowledge + informed opinions/beliefs
  - Combines multiple abilities in the best possible manner

"We are social beings"

Aristotle

MAS in the Web of Tomorrow

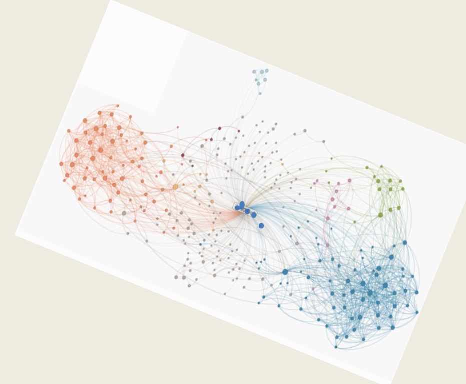

- Man-Agents collaboration
mediated by digital interaction technologies
  - combining human **and** machine intelligence
  - exploiting diverse large-scale collectives
  - …to solve complex problems

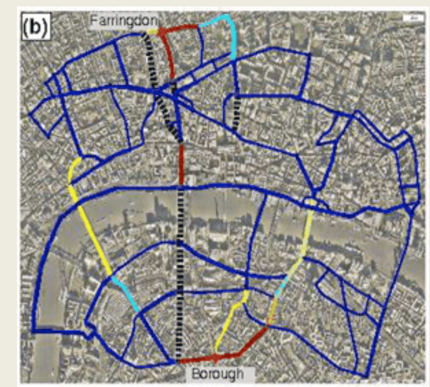

- Examples:
  - optimizing the transportation system of a city
  - managing patient treatment plans
  - search and rescue operations
  - coordinating communities of
    - activists;
    - Smart Grid consumers/producers

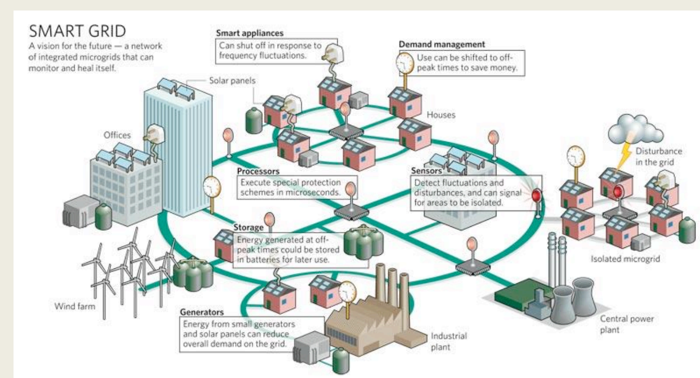

Thus, it's ok to not build the super-human, omnipotent intelligent agent

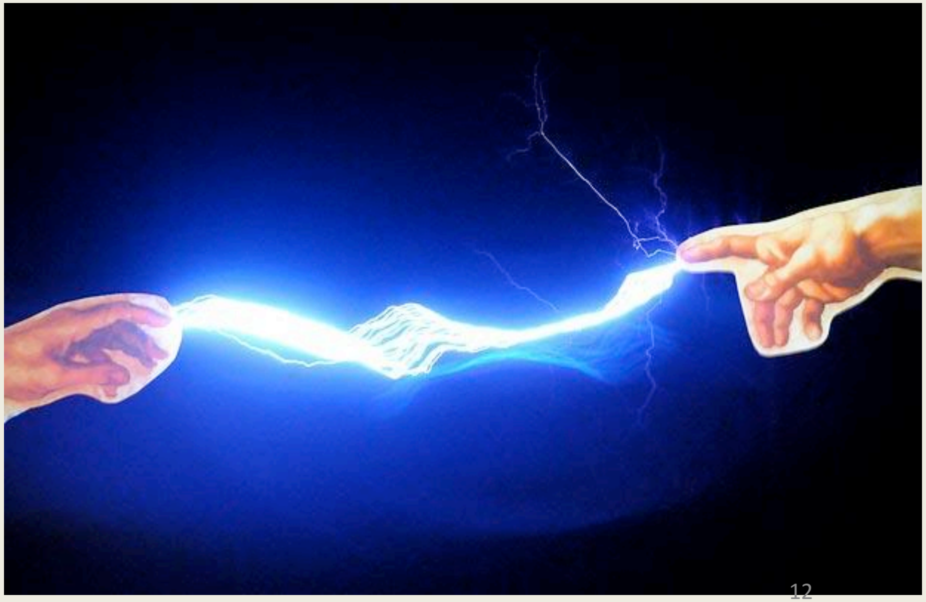

…as long as we get to get things done!…

- …to ease our everyday living…and even save lives!

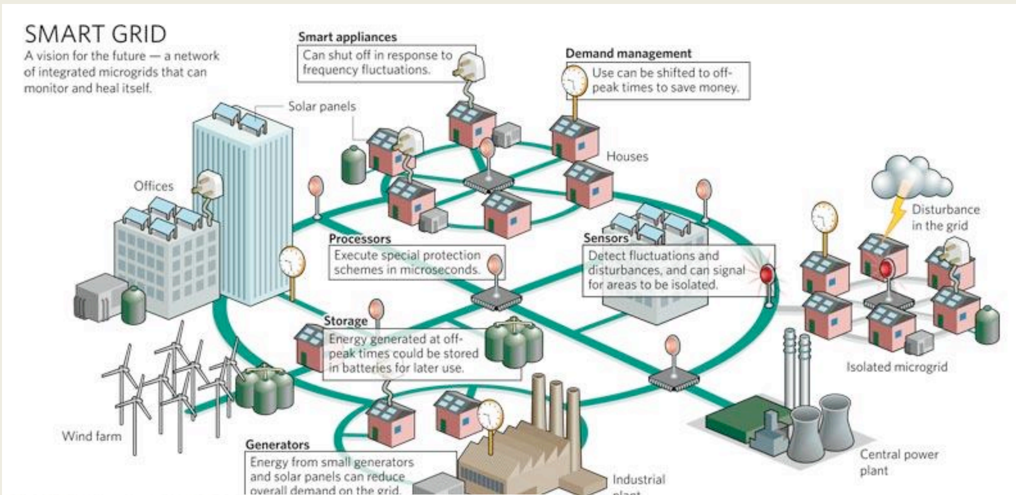

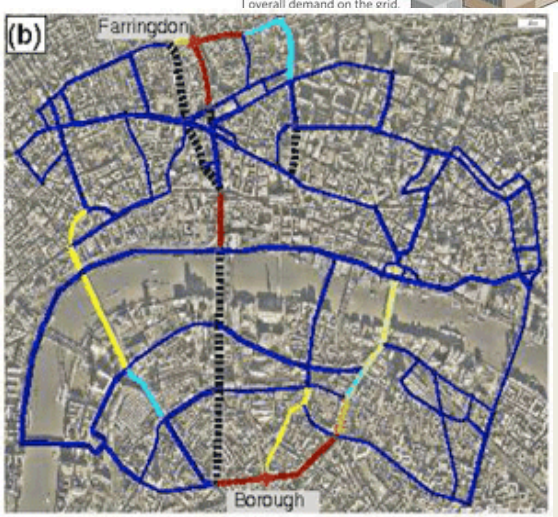

In a multi-agent world, you need to take decisions *and* act…

- …hopefully to your benefit, i.e., in a rational manner

Rational Choice Theory

- Finite set of actions A for a player,
  every a in A → outcomes
- At any given moment: choose some a
  from a subset of A
- How would a rational agent choose an action?

Example: Buying Oranges

- 3 varieties of oranges…from
  Sicily (S), Calabria (C), Puglia (P)

	S, P	P, C	S, P, C
Agent 1	S	C	C
Agent 2	P	P	S

- Who is rational; Who is (definitely) irrational?
- How do we see this using math?

Preferences

- Choices are based on preferences
- Actions/results: (a, b)
  - You either prefer a : (a > b), or b: (b > a)
- Consistency:
  If (a > b and b > c) then a > c

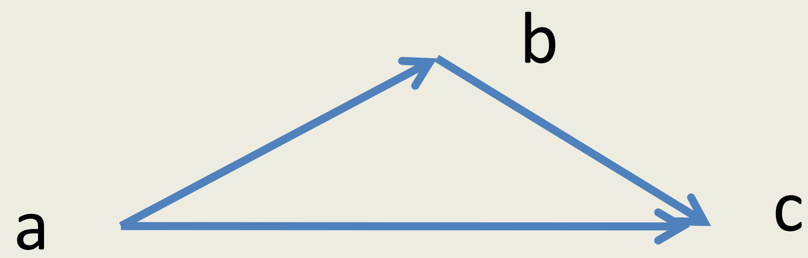

Consistent Preferences and Rationality

- **Observation:** If preferences are consistent, then for every set A there exists a in A so that a > b for every b in A \ {a}

- **Question:** given this, when will an agent be considered rational?

Consistent Preferences and Rationality

- **Observation:** If preferences are consistent, then for every set A there exists a in A so that a > b for every b in A \ {a}

- **Definition:** an agent is rational if she has consistent preferences $\succ$ and given an action set A the agent chooses a in A so that a $\succ$ b for every b in A \ {a}

Buying Oranges ++

- 3 varieties of oranges…from Sicily (S), Calabria (C), Puglia (P)

	S, P	P, C	S, P, C
Agent 1	S	C	C
Agent 2	P	P	S

- Agent 1's choices are compatible with the consistent preference relation $C > S > P$
- By contrast, agent 2's choices are incompatible with *any* consistent preferences
  - First choice suggests $P > S$, third choice suggests $S > P$

Weak Preferences

- The agent might be indifferent between two actions/outcomes
  $a \sim b$
- Non-strict/weak preferences
  - We write $a \succsim b$ if $a \succ b$ **or** $a \sim b$
- Consistency (or Transitivity): $a \succsim b$ and $b \succsim c$ implies $a \succsim c$
- Rationality: given A, the agent chooses a s.t. $a \succsim b$ for every b in $S \setminus \{a\}$
- *Claim*: with weak preferences allowed, we can no longer ascertain that an agent is irrational.

Buying Oranges ++

- 3 varieties of oranges…from
  Sicily (S), Calabria (C), Puglia (P)

	S, P	**P, C**	**S, P, C**
Agent 1	S	C	C
Agent 2	P	P	S

- Agent 2 might simply be indifferent among all choices

Preferences and Utilities

- $u: A \rightarrow \mathbf{R}$ maps an action to a (real) number (utility function)
- u represents $\succ$ if the following holds:

  $u(a) > u(b)$ *if-f* $a \succ b$

- Note: a given preferences' set can be represented by a multitude of utility functions

Buying oranges ++

- 3 varieties of oranges…from Sicily (S), Calabria (C), Puglia (P)

	S, P	P, C	S, P, C
Agent 1	S	C	C
Agent 2	P	P	S

- Agent 1’s choices, $C \succ S \succ P$, can be represented as
  $\{u(C) = 3, u(S) = 2, u(P) = 1\}$, or as
  $\{u(C) = 100, u(S) = 10, u(P) = 0\}$, or as
  $\{u(C) = 3572, u(S) = 9, u(P) = 5\}$ etc.

Preferences vs. Utilities

- With the addition of utilities, Rational Choice Theory is turned into Utility Theory...
  - **Advantages**?
- **1st** Preferences represented by providing : a list of *n* numbers instead of *n(n-1)/2* pairs
- **2nd** More "intuitive" representation, since preferences might be influenced or determined by "values":
  - The prices of the oranges might be what determines our preferences as to which ones to buy

Utilities++

- Where else does the need for utilities come from?
- …uncertainty / action stochasticity...
- ... Actions with uncertain outcomes can be seen as corresponding to lotteries wirh possible “prizes”

Lotteries

- A simple standard lottery L leads to prize A with probability *p*, and to B with *1-p*.

- More generally: $L = \{(p_1, X_1), (p_2, X_2), \ldots, (p_n, X_n)\}$
  - …and outcomes X_i might also be lotteries themselves!

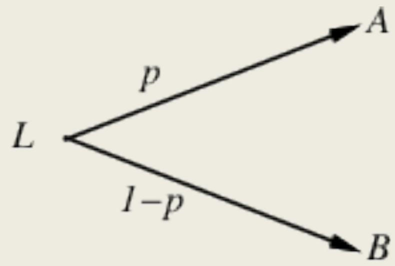

Lotteries++

- Given certain "intuitive" preference/ lotteries –related assumptions (e.g. consistency, monotonicity, etc), we can prove:

The Ramsey/ von Neumann & Morgenstern Utility Theorem

One can construct a real-valued utility function U which corresponds to the expected utility of a lottery, such that:

$U(L1) \geq U(L2)$ **if-f** $L1 \succsim L2$

where $U(L) = \Sigma\, p_i\, U(X_i)$

Lotteries & Utilities ++

- Ok, but...
- ... A lottery corresponds to a potential action with uncertain outcomes ... thus:
- The theorem leads to the Expected Utility Maximization Principle...

Maximum Expected Utility Principle

- Always choose the action with maximum expected utility
- Why;
- Because it corresponds to that lottery which is the most preferable among all lotteries!

Rationality (again…)

- Thus, *rational is the agent who acts according to the Expected Utility Maximization Principle*
- Therefore, yet another (**3rd**) advantage of using utilities instead of preferences: they can be used for decision making under uncertainty
- With the addition of probabilities/ expectations, Utility Theory is turned into Decision Theory...

Rational Choice /Utility Theory vs. Game Theory

- Rational Choice / Utility Theory: a single agent, preferences wrt actions
- Game Theory: many agents, preferences wrt strategy profiles (= vectors of actions of all players)

Strategic Games

- A strategic game consists of:
  - a set of players
  - for each player, an action set
  - for each player, preferences regarding the action profiles = vectors of all players' actions)
- Example:
  - Players A and B
  - Actions of player A: a1, a2 . Actions of B: b1, b2
  - Preferences A: (a1, b1) **>** (a2, b2) **>** (a1, b2) **>** (a2, b1)

Prisoner's Dilemma

- Two guys arrested
- Enough evidence for a 1 year conviction for each, but only indications of a serious crime (4 years)
- *One* confesses → walks free, the other gets 4 years
- *Both* confess → 3 years in jail each

Normal Form (or Matrix) Representation

P1 \ P2	“omertà”	“snitching”
“omertà”	(-1,-1)	(-4, 0)
“snitching”	(0, -4)	(-3, -3)

- **By convention:** a pair (x, y) where row i meets column j means that the row player who chooses i gets x, while the “column” player who chooses j gets y

Prisoner's Dilemma: Preferences

P1 \ P2	Cooperate	Defect
Cooperate	(-1,-1)	(-4, 0)
Defect	(0, -4)	(-3, -3)

P1(row): (D, C) > (C, C) > (D, D) > (C, D)

P2 (column): (C, D) > (C, C) > (D, D) > (D, C)

Interlude:
Prisoner's Dilemma: The dilemma

P1 \ P2	Cooperate	Defect
Cooperate	(-1,-1)	(-4, 0)
Defect	(0, -4)	(-3, -3)

P1(row): (D, C)

P2 (column):(C, D)

What would you do?

Is there always a dilemma?

- what if they could communicate?
  - What if they could “sign a contract”?
  - ➢ *cooperative game theory*
- what if the game was to be repeated?
  - ➢ *for a finite or for an infinite number of rounds?*

Common Project

P1 \ P2	Lavorare	Giocare
Lavorare	(3, 3)	(-5, 5)
Giocare	(5, -5)	(0, 0)

Two students per project. *At least one* works→project suceeds. Project succeds-→each student happy (+5), but each student would rather not work (lavorare→ -2), but (s)he does not want to be exploited (if lavorare while the other giocare → -8 for him/her)

Common Project vs. Prisoner's Dilemma

	C	D
C	(3, 3)	(-5, 5)
D	(5, -5)	(0, 0)

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

- "Common project", column player preferences. (D, C) > (C, C) > (D, D) > (C, D)
- Two games are equivalent if they imply the same preferences wrt choices profiles

Hawk – Dove (or “Chicken”)

P2 P1	Dove	Hawk
Dove	(3, 3)	(1, 4)
Hawk	(4, 1)	(0, 0)

Is this game equivalent to the PD?

Mozart or Mahler?

P2 / P1	Mozart	Mahler
Mozart	(2, 2)	(0, 0)
Mahler	(0, 0)	(1, 1)

No coordination → they both lose

Not all coordinated outcomes are equally good

Identical Interest or Coordination Games

All players have the same payoff function

…which is…

Possibly stochastic

	A	B
A	X, X	C, C
B	D,D	M, M

Battle of the Sexes

P1 P2	Teatro	Calcio
Teatro	(2, 1)	(0, 0)
Calcio	(0, 0)	(1, 2)

Will we ever meet? And where?

Matching Pennies

P1 \ P2	Testa	Croce
Testa	(1, -1)	(-1, 1)
Croce	(-1, 1)	(1, -1)

TT o CC → P1 wins; otherwise P2.

Matching Pennies

P1 \ P2	Testa	Croce
Testa	(1, -1)	(-1, 1)
Croce	(-1, 1)	(1, -1)

A strictly competitive game

Zero-sum games: $u_i(a) + u_j(a) = 0$

A beautiful mind

Characterized rational play for
non-zero sum strategic games

John Forbes Nash, Jr.

Nobel Prize for Economics, 1994
4 Oscars, 2002

Nash equilibrium (informally)

Nash equilibrium states: no player can profit by **unilaterally** switching her action, given other players' actions.

Nash equilibrium (informally)

...that is, assuming rational agent i plays a_i, rational agent j has no (rational) choice but to play a_j

...

and, assuming rational agent j plays a_j, rational agent i has no (rational) choice but to play a_i

Nash Equilibrium (formally)

- Given a strategic game with *n* players…
  - player i chooses action a_i from A_i , resulting to utility u_i
- …an action profile $\underline{\mathbf{a}} = (a_1,.. a_i .., a_n)=(a_i , \underline{\mathbf{a}}_{-i})$,
  is a Nash equilibrium if:

  for **every** player $i = 1, ..., n$, we have

  $u_i (\underline{\mathbf{a}}) \geq u_i (\underline{\mathbf{a}}_{-i} , a')$ **for all** a' in A_i

Nash equilibrium visually

(,)	(,)	$(x_1, \quad)$	(,)	(,)
(,)	(,)	$(x_2, \quad)$	(,)	(,)
(,)	(,)	$(x_3, \quad)$	(,)	(,)
$(\quad , y_1)$	$(\quad , y_2)$	(X, Y)	$(\quad , y_4)$	$(\quad , y_5)$
(,)	(,)	$(x_5, \quad)$	(,)	(,)

X : at least as big as x_i,

Y : at least as big as y_j

Nash equilibrium ++

- Obviously: No player has an incentive to unilaterally move from a Nash equilibrium state
- **Sfortunamente:**
  - Solving for Nash is impossible in most realistic multiagent settings [Daskalakis 2008]
  - Many games do not have a Nash equilibrium (in pure strategies ...)
  - Many games have more than one Nash equilibrium
    - Not all NE are created equal / as desirable!
    - Which one to choose?
- **Attenzione:** the NE does not imply that alternatives are worse, but simply that they are not better than NE

NE Examples

Battle of the Sexes

T: Teatro
C: Calcio

P1 \ P2	T	C
T	(2, 1)	(0, 0)
C	(0, 0)	(1, 2)

- Is (T, C) Nash equilibrium? Why yes/not;
- (T, T) and (C, C): Nash equilibria
- P1 ’s most preferred action profile is **not** (C, C), but if there, no incentive to switch: if he switches to *Teatro*, they play (T, C) and he gets 0

Battle of the Cinefils

	Film 1	Film 2
Film 1	(3, 3)	(0, 0)
Film 2	(0, 0)	(1, 1)

- (F1, F1) : Nash equilibrium
- ... (F2, F2) also! Which one would you prefer?
- (F1, F1) : higher social welfare (sum of players' utilities)

…another battle of cinefils

	Horror	Social
Horror	(1, 9)	(0, 0)
Social	(0, 0)	(4, 4)

- (H, H) και (S, S): Nash equilibria
- Which one is more fair in your opinion?
  - (S, S) : more fair

Prisoners' Dilemma: Nash Equilibrium

P1 \ P2	C	D
C	(-1, -1)	(-4, 0)
D	(0, -4)	(-3, -3)

(D, D) is a Nash equilibrium

- $u_1(C, D) \leq u_1(D, D)$
- $u_2(D, C) \leq u_2(D, D)$

- (C, C) not a Nash equilibrium:

  P1 can switch to D → $u_1(D, C) > u_1(C, C)$

- (C, D) not a Nash equilibrium:

  P1 can switch to D → $u_1(D, D) > u_1(C, D)$

  Similarly, P2 can switch from (C,C), (D,C)

Best Response Functions

- Given opponents' action vector $\underline{\mathbf{a}}_{-i}$, player i has some actions that maximize her utility
- Best response "function":
  $B_i(\underline{\mathbf{a}}_{-i}) =$
  $\{a \text{ in } A_i \mid u_i(\underline{\mathbf{a}}_{-i}, a) \geq u_i(\underline{\mathbf{a}}_{-i}, a') \text{ for all } a' \text{ in } A_i\}$
- $B_i(\underline{\mathbf{a}}_{-i})$ is "multi-valued"
- If $|B_i(\underline{\mathbf{a}}_{-i})| = 1$ for all $\underline{\mathbf{a}}_{-i}$, we 'll write $b_i(\underline{\mathbf{a}}_{-i})$ for i's unique "best response" to $\underline{\mathbf{a}}_{-i}$

Example

	L	C	R
T	(2*, 5*)	(3 , 3)	(6*, 3)
M	(2*, 7*)	(4 , 5)	(2 , 7*)
B	(1 , 4*)	(5*, 4*)	(2 , 1)

- $B_1(L) = \{T, M\}$
- $B_1(C) = \{B\}$
- $B_1(R) = \{T\}$

- $B_2(T) = \{L\}$
- $B_2(M) = \{L, R\}$
- $B_2(B) = \{L, C\}$

Best Responses και Nash Equilibria

$\underline{\mathbf{a}} = (a_1, \ldots.., a_n)$ is a Nash equilibrium if-f
$u_i(\underline{\mathbf{a}}) \geq u_i(\underline{\mathbf{a}}_{-i}, a')$
for every i and for all a' in A_i

- from Best Responses definition, equivalently:

$\underline{\mathbf{a}} = (a_1, \ldots.., a_n)$ is a Nash equilibrium if-f
a_i belongs to $B_i(\underline{\mathbf{a}}_{-i})$ for every i

Therefore…

	L	C	R
T	(2*, 5*)	(3 , 3)	(6*, 3)
M	(2*, 7*)	(4 , 5)	(2 , 7*)
B	(1 , 4*)	(5*, 4*)	(2 , 1)

- $B_1(L) = \{T, M\}$, $B_1(C) = \{B\}$, $B_1(R) = \{T\}$
- $B_2(T) = \{L\}$, $B_2(M) = \{L, R\}$, $B_2(B) = \{L, C\}$
- **{T, L}, {M, L} and {B, C} are Nash equilibria**

Dominance

- Some times, some strategies are clearly better than others…

Strict Dominance

P1 \ P2	T	C
T	(5, 1)	(4, 4)
C	(0, 0)	(1, 5)

- P1 has a strong preference for Teatro
  - prefers T instead of C regardless of P2's choices
- P2 really prefers Calcio
  - prefers C instead of T no matter what P2 chooses

Strict Dominance

Action a strictly dominates an action b of i, if i prefers every outcome of playing a when compared to any outcome of playing b

Strict Dominance: Definition

- An action a of i
  strictly dominates an action b (of that same player) if
  $$u_i(\underline{\mathbf{a}}_{-i}, a) > u_i(\underline{\mathbf{a}}_{-i}, b)$$
  for any profile $\underline{\mathbf{a}}_{-i}$ of other players' actions.
- If b is *strictly dominated* by a, player i *cannot possibly play* b in a Nash equilibrium
  - …since she can unilaterally deviate profitably to a

Use dominance to find Nash

- A rational agent will never choose a strictly dominated strategy
- Thus the strictly dominated strategies can be eliminated from the game

Strict Dominance and NE

Π1 \ Π2	T	C
T	(5, 1)	(4, 4)
C	(0, 0)	(1, 5)

- For the row player, T strictly dominates C
- For the column player, C strictly dominates T
- Thus, (T, C) is the (unique) NE

Prisoner's Dilemma

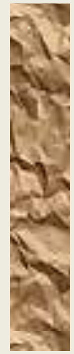

	C	D
C	(-1, -1)	(-4, 0)
D	(0, -4)	(-3, -3)

For both players,
D strictly dominates C

(D, D) is the unique NE

Iterated elimination of strictly dominated strategies

- By this elimination process, no NE is lost
- Why? "I am rational, and I know you are rational, and I know you know I am rational, and you know that I know these, ...thus we will never use these strictly dominated strategies"
- Thus, we use this process to find NE via simplifying the original game

Mixed strategy NE (Nash equilibrium in mixed strategies)

- Mixed strategies:
  - Choose an action based on a probability distribution
  - pure NE: subset of a game's mixed NE

Mixed strategy NE (Nash equilibrium in mixed strategies)

- Mixed strategies:
  - play a_1 with probability p_1
  - play a_2 with probability p_2
  - . . .
  - play a_k with probability p_k

  s.t. $p_1 + p_2 + \quad + p_k = 1$

- Nash (1951): all finite games have at least one NE in mixed strategies

Computing mixed NE

Matching pennies

P1 \ P2	Corona	Lettere
Corona	(1, -1)	(-1, 1)
Lettere	(-1, 1)	(1, -1)

Computing mixed NE ++

- P1 : C with probability *p*
- P2 : C with probability *q*

If P2 plays mixed,
he must be indifferent between C / L
(or he would have played a pure strategy)

$$u_2(C) = u_2(L)$$
$$-p+(1-p) = p-(1-p)$$
$$p = 1/2$$

Computing mixed NE ++

- Similarly, P1 indifferent between C/ L ..
- ...thus, p=1/2, and q=1/2

Computing mixed NE ++

- If all players are indifferent among using pure strategies, probabilities are defined for others' mixes so they each play a best response to others' strategies
  - ...if any other mix was used, opponents could have have chosen a (pure) strategy that would have improved her payoff
  - ...thus agents play a Nash equilibrium (vector of best responses)

Computing mixed NE: The stag hunt game

P1 \ P2	Deer	Rabbit
Deer	(4, 4)	(1, 3)
Rabbit	(3, 1)	(3, 3)

Assume P1 plays "Deer" with probability p, P2 plays "Deer" with probability q.

This holds since 2 plays a BR to 1's mixed strategy… *u2("Deer") = u2("Rabbit")*

$$4p+1(1-p) = 3p + 3(1-p)$$

...a fact that fixes a p for 1's BR, given which 1 is indifferent: $p = 2/3$

u1("Deer") = u1("Rabbit")

$$4q+(1-q) = 3$$

$$q = 2/3$$

The stag hunt: (expected) utilities at mixed NE

P1 \ P2	Deer	Rabbit
Deer	(4, 4)	(1, 3)
Rabbit	(3, 1)	(3, 3)

$$EU[1] = 4pq+p(1-q)1+3(1-p)q + 3(1-p)(1-q)$$
$$= 4x2/3x2/3 + 2/3x1/3+3x1/3x2/3+3x1/3x1/3$$
$$= 16/9 + 2/9 + 6/9 + 3/9 = 3$$

$$EU[2] = 4pq+3p(1-q)+(1-p)q + 3(1-p)(1-q)$$
$$= 4x2/3x2/3 + 3x2/3x1/3+1/3x2/3+3x1/3x1/3$$
$$= 16/9 + 6/9 + 2/9 + 3/9 = 3$$

The stag hunt: (expected) utilities at mixed NE

P1 \ P2	Deer	Rabbit
Deer	(4, 4)	(1, 3)
Rabbit	(3, 1)	(3, 3)

- **In NE all players play best responses. If I tell you that any mixed BR has an EU that is equal to that of the pure strategies in the mix (it's true!), can you think of an easier way to calculate these quantities?**

EU[1] = 4pq+p(1-q)1+3(1-p)q + 3(1-p)(1-q)

= 4x2/3x2/3 + 2/3x1/3+3x1/3x2/3+3x1/3x1/3

= 16/9 + 2/9 + 6/9 + 3/9 = 3

EU[2] = 4pq+3p(1-q)+(1-p)q + 3(1-p)(1-q)

= 4x2/3x2/3 + 3x2/3x1/3+1/3x2/3+3x1/3x1/3

= 16/9 + 6/9 + 2/9 + 3/9 = 3

Computing mixed NE ++

- Note: causing *“indifference” to others is not the agents’ goal*
  - They just try to maximize their payoff (playing best responses)...
  - ...but ”indifference among pure stategies” is a *requirement* for mixed NE

Justifying mixed strategy NE

- attempt to insert randomization in play
- attempt to include a statistical reasoning about others behaviour
  - NE as a steady state in a stochastic game
- models dependence on “outside”/ unknown factors
- models player strategies with unstable or varying preferences

NE is not always the appropriate rationality concept

Finite games in extensive form: Subgame-Perfect Equilibrium

(example: negotiations with multiple rounds)

Extensive form games: “empty” threats

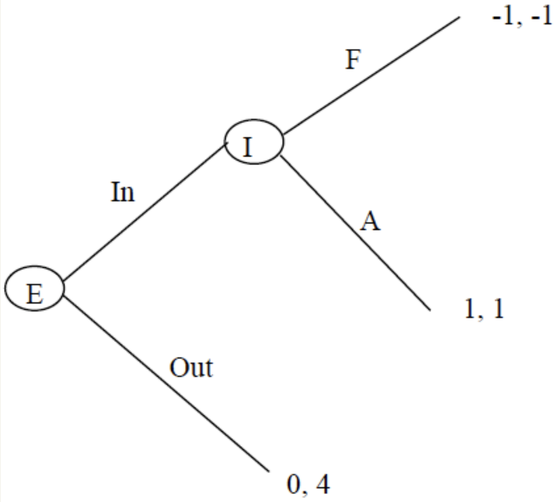

Pure Strategies: $S_E = \{In, Out\}$ and $S_I = \{F, A\}$.

Normal Form representation:

	F	A
In	-1, -1	$\boxed{1, 1}$
Out	$\boxed{0, 4}$	0, 4

Both (pure) NE, but (Out, F) corresponds to a non-credible threat: “I” will never choose “Fight” if E plays “In”

Subgames in extensive form games with perfect information

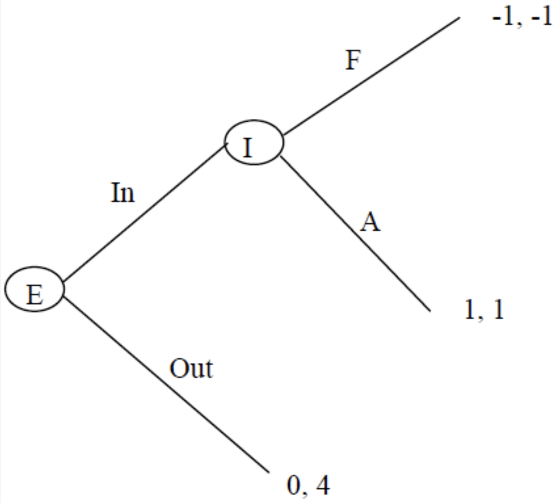

Subgame: A game subset that includes an "initial" node which the game has reached, and all nodes-children of that node, as well as their descendents.

The game here has two subgames!

With "imperfect information", the definition is a bit different and includes "information sets".

Subgame-Perfect Equilibrium

Subgame-perfect equilibrium:
Nash equilibrium in every subgame!

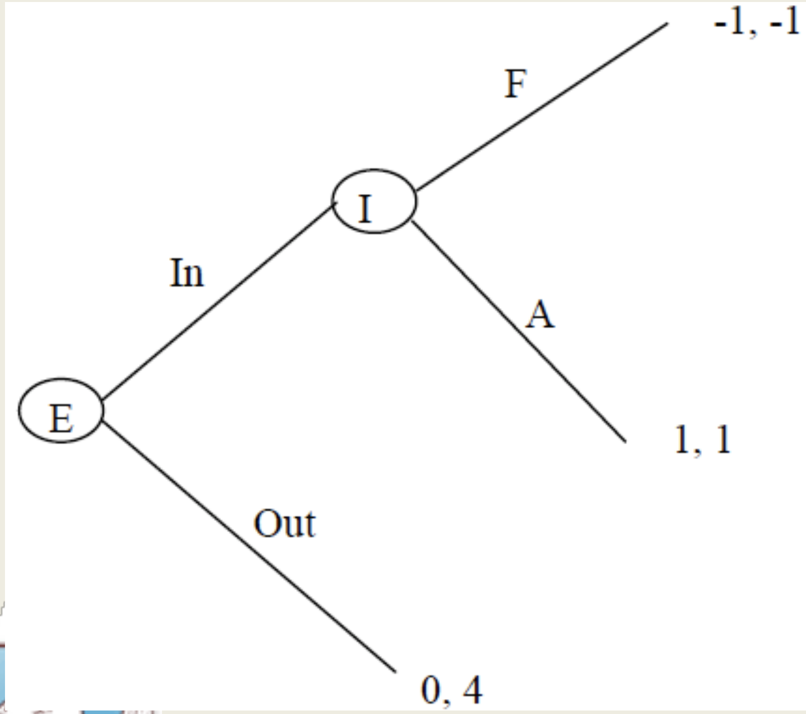

(Out, F): equilibrium in the ”big”subgame (whole game), but not the subgame after “In”
With backward induction: given “Accommodate”, E cannot but play “In” in the big subgame!

Careful: in extensive form games, we need complete plans of action

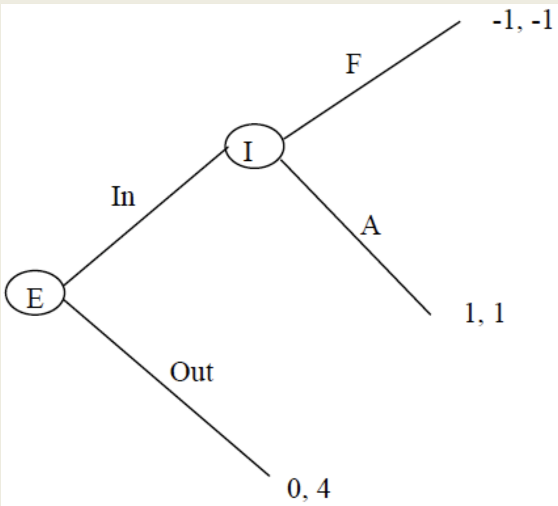

Pure Strategies: $S_E = \{In, Out\}$ and $S_I = \{F, A\}$.

Normal Form representation:

	F	A
In	-1, -1	1, 1
Out	0, 4	0, 4

F corresponds to "F if E plays In", and A to "A if E plays In"

Dividere l’anguria

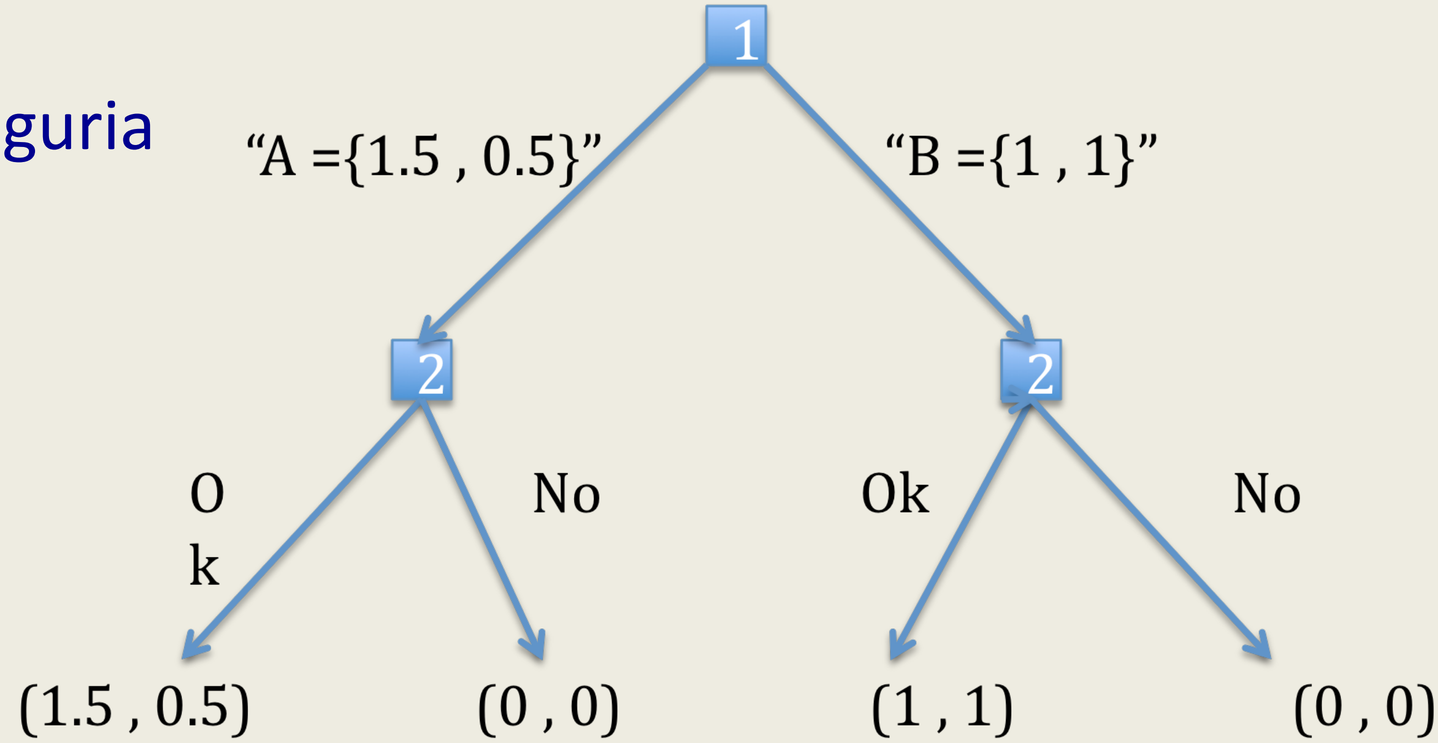

You have to have a complete action plan!

Player 2

Player 1

	Ok|A and Ok|B	Ok|A and No|B	No|A and Ok|B	No|A and No|B
A	1.5, 0.5	1.5, 0.5	0, 0	0,0
B	1, 1	0, 0	1, 1	0, 0

Backward induction

starting by the last round (of a game in extensive form), define inductively the sequence of optimal actions

Suggested Readings

- K. Leyton-Brown & Y. Shoham: “Essentials of Game theory: A concise multidisciplinary introduction”: https://doi.org/10.2200/S00108ED1V01Y200802AIM003

Additional Slides

Repeated Prisoner's Dilemma

- If you know you are going to face the dilemma over and over again, then snitching is probably not a good idea...:

Infinitely repeated PD: cooperation is the rational choice

...but, in the real world, nothing lasts for ever...☹

…thus …with finite horizon…

- If known number of repetitions *n*
- …then at n round you have an incentive to defect...
- Both “snitch” on the other at n…
- …making *n-1* the last “real” round...
  - Thus, incentive to snitch at *n-1*
- ...with backward induction we can prove that

In finitely repeated PD:

to defect is the rational choice (again)

…unless you are not sure of the value of *n* ... ☺

Repeated Games

- Repeated execution of a stage game:

	s	t
q	ι,κ	λ,μ
p	ο, π	ν, ξ

"stage game"

Repeat

- In each repetition, players move simultaneously, observe others' moves after the round is over
- Players' payoff functions are additive

Taking Decisions over Time: Stochastic or Markov games

- Beyond repeated games:

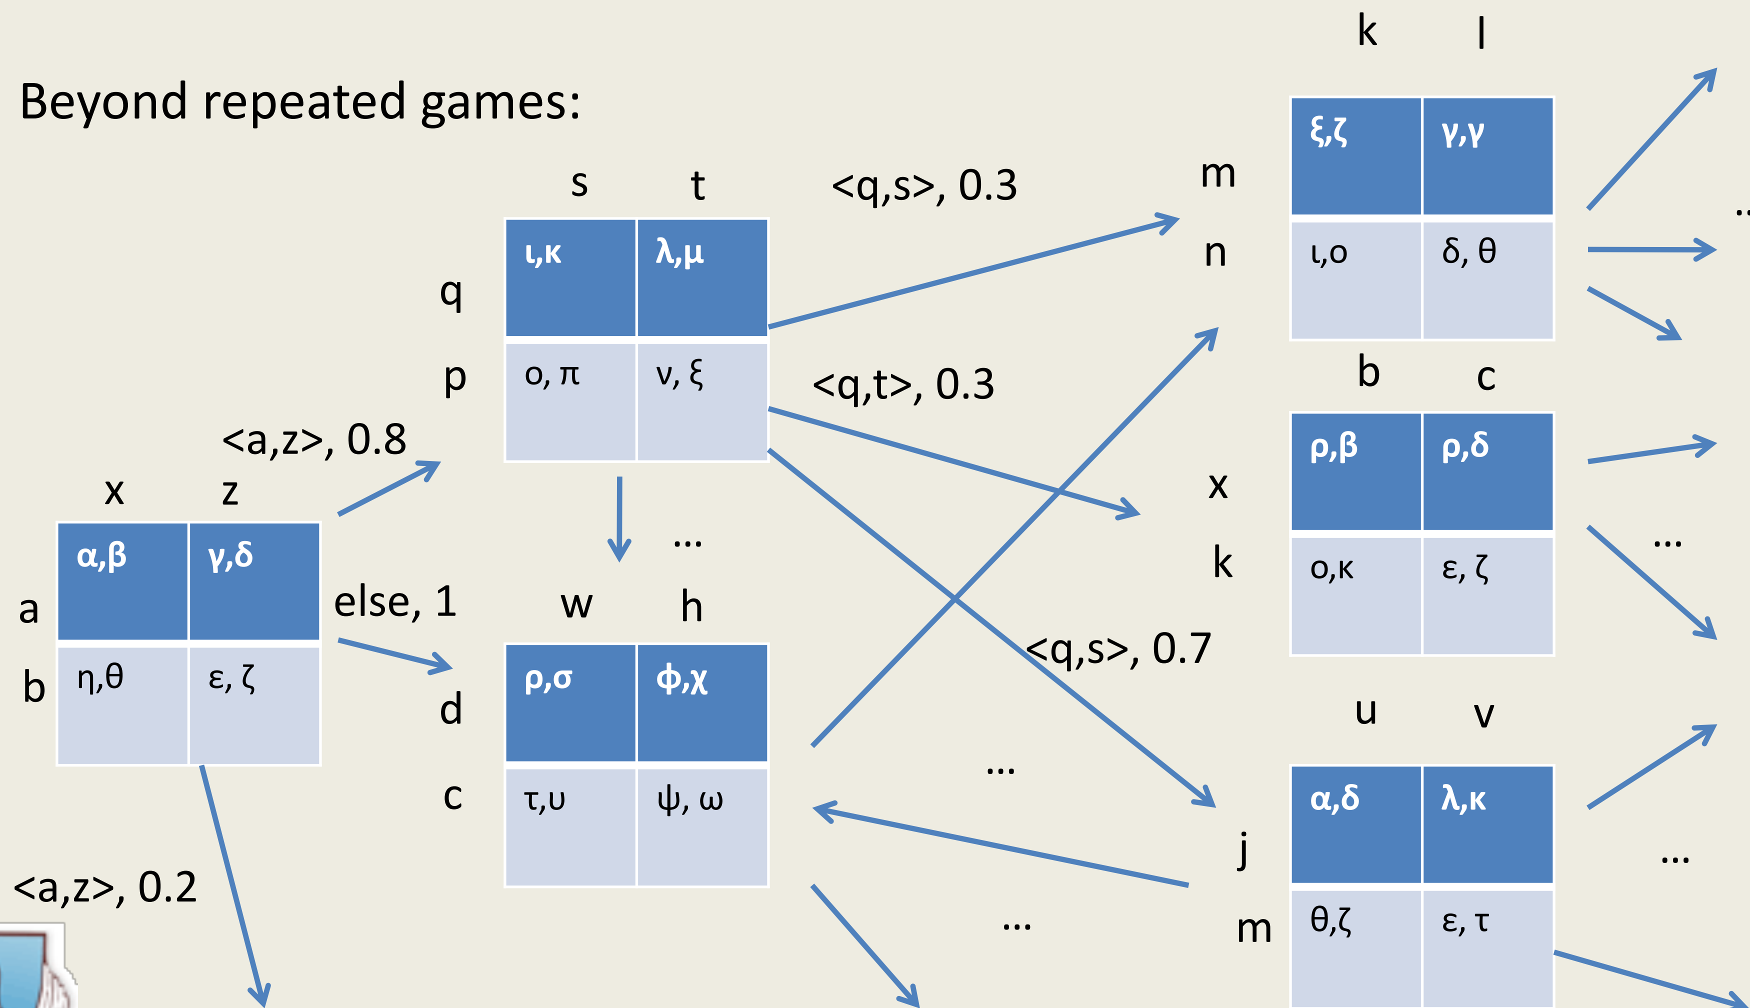

Symmetric games

- Can players swap IDs without a need to change strategies (swapped players' strategies result to same payoffs)?
- Then the games are symmetric
- Lavoro: Check to see which of the aforementioned games are symmetric!

(a, a)	(b, c)
(c, b)	(d, d)

Two algorithms to identify NE in 2-player games

1. For every cell, check if NE, that is check whether
   (1) P1 can benefit by switching
   (2) P1 can benefit by switching

2. (1) for every P2's action, compute P1's best responses
   (2) for every P1's action, compute P2's best responses
   (3) display the cells where best responses meet

- Which one is the most efficient computationally?

Algorithm 1

- 2-player n-by-n game (every player has *n* actions)
- For every cell, check if NE, that is check whether:
  - (1) P1 can benefit by switching: *n* checks
    (2) P2 can benefit by switching: *n* checks
  - n^2 cells, thus $(n+n) \times n^2 = 2n^3$ computational cost

(,)	(? ,)	(,)
(, ?)	(x , y)	(, ?)
(,)	(? ,)	(,)

Algorithm 2

1. for every action (column) of P2, compute the best responses of P1:
   - Scan column to identify P1's max payoff at every column
   - Scan column again to mark cells max payoff
   - Each scan: requires n checks, thus 2n x n checks in total for this part
2. same for P1's actions
3. Scan matrix to identify

Cells with two labels: n^2

In total:

- $2n^2 + 2n^2 + n^2 =$ $5n^2$ checks
- for large *n*, $5n^2 < 2n^3$

(?*,?*)	(? ,?)	(?*, ?)
(? ,?)	(?*,?*)	(? , ?)
(? ,?)	(? ,?*)	(? , ?)

Best Response Functions: Continuous Action Spaces

- Two programmers work together in a project
- Workload/effort of each player: a number in [0, 1]
- If 1 allocates x units of effort, and
  2 allocates y units of effort,
  and 1's utility is x(c+y-x),
  while 2's is y(c+x-y)
- best response of 1 to y: (c+y)/2
- best response of 2 to x: (c+x)/2

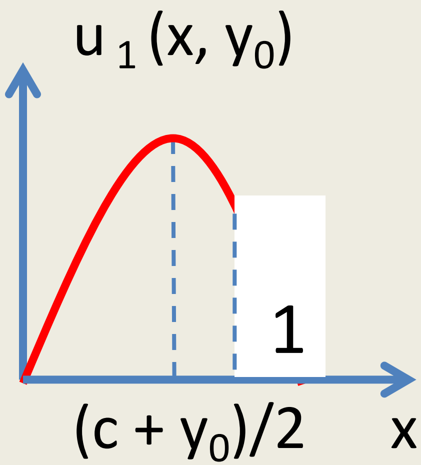

Common project in continuous spaces, algebraic solution

- best response of 1 to y: (c+y)/2
- best response of 2 to x: (c+x)/2
- (x, y) is Nash equilibrium if
  - x best response of player 1 στο y
  - y best response of player 2 στο x
- x = (c+y)/2, y = (c+x)/2
- 2y = c+(c+y)/2   =>   4y = 3c+y
- Thus: y = c, x = c

Common project in continuous spaces, algebraic solution ++

- That is, we start by differentiating the utility functions (to define best response functions, which return actions that maximize utility as a function of opponent actions), and then solve simultaneously the equations that describe the best responses of players to opponent actions...