COOPERACIÓN

Cooperation

2009

Master Course

13/12/09

Self-interested Agents

- What does it means to that an agent is self-interested?
  - Not that they want to harm other agents
  - Not that they only care about things that benefit them
  - That the agent has its own description of states of the world that it like, and its actions are motivated by this description

Utility Theory

- Quantifies degree of preference across alternatives
- Understand the impact of uncertainty on these preferences
- **Utility function**: a mapping from states of the world to real numbers, indicating agent's level of *happiness* with that state of the world
- Decision-theoretic rationality: take actions to maximize expected utility

Preferences over Outcomes

- If o_1 and o_2 are outcomes
  - $o_1 \geq o_2$ means o_1 is at least as desirable as o_2.
  - read this as "the agent *weakly prefers* o_1 to o_2"
  - $o_1 \sim o_2$ means $o_1 \geq o_2$ and $o_2 \geq o_1$
  - read this as "the agent is indifferent between o_1 and o_2"
  - $o_1 > o_2$ means $o_1 \geq o_2$ and $o_2 \neg \geq o_1$
  - read this as "the agent strictly prefers o_1 to o_2"

Lotteries

- An agent may not know the outcomes of his actions, but may instead only have a probability distribution over the outcomes.
- A Lottery is a probability over outcomes. It is defined as:
  - $[p_1: o_1, p_2: o_2, p_1: o_1\dots.\ p_k: o_k,]$

  Where the o_i are outcomes and $p_i > 0$ such that

$$\sum_i p_i = 1$$

- The lottery specifies that outcome o_i occurs with probability p_i.
- Lotteries are outcomes

Preference Axioms (1)

- Completeness: A preference relationship must be defined every pair of outcomes:

$$\forall o_1 \forall o_2 \; o_1 \geq o_2 \textbf{ or } o_2 \geq o_1$$

dmerida@lsi.upc.edu & ia@lsi.upc.edu

COOPERATION

Preference Axioms (2)

- Transitivity: Preference must be transitive:

$$\textbf{if } o_1 \geq o_2 \textbf{ and } o_2 \geq o_3 \textbf{ then } o_1 \geq o_3$$

- This makes good sense otherwise
- **if** $o_1 \geq o_2$ **and** $o_2 \geq o_3$ **then** $o_3 \geq o_1$

- An agent should be prepared to pay some amount to swap between an outcome he prefers *less* and an outcome the prefer ***more***

dmerida@lsi.upc.edu & ia@lsi.upc.edu

COOPERATION

Preference Axioms (3)

- Monotonicity: Preferences should preserve order

$$\textbf{if } o_1 \geq o_2 \textbf{ and } p \geq q \textbf{ then}$$
$$[p:o_1, 1-p:o_2] > [q:o_1, 1-q:o_2]$$

- An agent prefers a larger chance of getting a better outcome to a smaller chance

dmerida@lsi.upc.edu & ia@lsi.upc.edu

Preference Axioms (4)

- Decomposability: *No fun in game*

$$\textbf{if } \forall o_1 \in O,\ P_{\ell}(o_1) = P_{\ell}(o_2)$$
$$\textbf{Then } \ell_1 \sim \ell_2$$

- Where $P_{\ell}(o_i)$ denotes the probability that outcome o_i is selected by lottery ℓ.

COOPERATION

dmerida@lsi.upc.edu & ia@lsi.upc.edu

COOPERATION

Preference Axioms (5)

- Continuity: suppose that $o_1 > o_2$ and $o_2 > o_3$, then there exists a $p \in [0,1]$ such that $o_2 \sim [p:o_1, 1- p:o_3]$.

dmerida@lsi.upc.edu & ia@lsi.upc.edu

Preference Axioms (6)

- Substitutability: If $o_1 \sim o_2$ then for all sequences of one or more outcomes $o_3 \ldots o_k$, and sets of probabilities $p, p_3, \ldots p_k$ for which $p + \Sigma_{(i=3,k)} p_i = 1$

$$[p{:}o_1, p{:}o_3, \ldots p_k{:}o_k] \sim [p{:}o_2, p{:}o_3, \ldots p_k{:}o_k]$$

COOPERATION

Preference and utility functions

- Theorem (Von Neumann and Morgenstern, 1944)

*If an agent's preference relation satisfies the axioms of Completeness, Transitivity, Decomposability, Substitutability, Monotonicity and Continuity **then** there exists a function u:O→[0,1] with the following properties:*

1. $u(o_1) \geq u(o_2)$ iff the agent prefers o_1 to o_2; and
2. when faced about uncertainty about which outcomes he will receive, the agent prefers outcomes that maximize the expected value of u.

dmerida@lsi.upc.edu & ia@lsi.upc.edu

GAME THEORY

2009

dmerida@lsi.upc.edu &
ia@lsi.upc.edu

Game Theory-Basic ideas

- **Game Theory**: Is a theory about the agents' **rational** behaviour in **interaction** problems of a **group** of agents showing a **strategic behaviour**.
  - Expressed in mathematical terms
  - Plays an important role in the today's economy

dmerida@lsi.upc.edu & ia@lsi.upc.edu

Rationality

- One of the most common assumptions made in game theory (along with common knowledge of rationality). In its mildest form, **rationality implies that every player is motivated by maximizing his own payoff**.
- In a stricter sense, it implies that every player *always* maximizes his utility, thus being able to *perfectly* calculate the probabilistic result of every action

Common Knowledge

- An item of information in a game is common knowledge if all of the players know it (it is mutual knowledge) and all of the players know that all other players know it and all other players know that all other players know that all other players know it, and so on. This is much more than simply saying that something is known by all, but also implies that the fact that it is known is also known by all, etc.
- Consider a simple example of two allied armies situated on opposite hilltops waiting to attack their foe. Neither commander will attack unless he is sure that the other will attack at exactly the same time. The first commander sends a messenger to the other hilltop with the message "I plan to attack in the morning." The messenger's journey is perilous and he may die on the way to delivering the message. If he gets to the other hilltop and informs the other commander - can we be certain that both will attack in the morning? Note that both commanders now *know* the message, but the first cannot be sure that the second got the message.
- Thus, **common knowledge implies not only that both know some piece of information, but can also be absolutely confident that the rest no it, and that the rest know that we know it, and so on**.

Mutual Knowledge

- Something in a game is *Mutual Knowledge* if all players know it.
- A seemingly simple concept, mutual knowledge is insufficient to analyze most games, since it is not clear from this assumption alone what people think others know. I might know that X is true, but my actions may depend on whether or not other players know that I know X.
- A common additional assumption is that the facts of the game are not only *mutual knowledge* but also *common knowledge*.
- This implies that we all know X, and we all know that everyone else knows X, and we all know that everyone knows that everyone else knows X, and so on.

Game Theory-Basic ideas

Game Theory: The games are well-defined mathematical objects.

A game consists of a set of players, a set of moves (or strategies) available to those players, and a specification of payoffs for each combination of strategies.

Most cooperative games are presented in the characteristic function form, while the extensive and the normal forms are used to define non-cooperative games.

- Group (>1 players)
- Interaction (Do they afect between them?)
- Estrategy (All know)
- Rational (Choice the best option)*

Game Theory-Basic ideas

- Game theory allows to model the mathematics of the interest in conflict.
- Its origin is the interest of modelling *games* as poker or chess but not the roulette.
- Game theory was born upon John von Neumann's interest in modelling poker (1940).

dmerida@lsi.upc.edu & ia@lsi.upc.edu

COOPERATION

Game Theory-Basic ideas

- Two central concepts in Game Theory are: payoffs and strategy
- An strategy is a program for a player: a sequence of actions
- When a game ends, each player gets a payoff.
  - **A payoff maybe positive or negative**

dmerida@lsi.upc.edu & ia@lsi.upc.edu

COOPERATION

Strategy

- ***A strategy defines a set of moves or actions a player will follow in a given game***.
- A strategy must be complete, defining an action in every contingency, including those that may not be attainable in equilibrium.
- For example, a strategy for the game of checkers would define a player's move at every possible position attainable during a game. Such moves may be random, in the case of mixed strategies.

COOPERATION

Payoff

- In any game, payoffs are numbers which represent the motivations of players.

- Payoffs may represent *profit*, *quantity*, *utility*, or other continuous measures (cardinal payoffs), or may simply rank the desirability of outcomes (ordinal payoffs).

- In all cases, the payoffs must reflect the motivations of the particular player.

Game Theory-Basic ideas

- According with Rubenstein: Game Theory is the study of the considerations that **may or may not** take in to account previous strategic situations
- Links with IA: Design models of strategic behaviours. **Players are rational.**

Game Theory

- A game, in theory game, can be defined as:
  - A game consists of a set of players, a set of moves (or strategies) available to those players, and a specification of payoffs for each combination of strategies.
  - This payoff depends of the players' decisions and, possibly, on *good luck*.

dmerida@lsi.upc.edu & ia@lsi.upc.edu

COOPERATION

GAME

- The interaction among rational, mutually aware players, where the decisions of some players impacts the payoffs of others.
- A game is described by its players, each player's strategies, and the resulting payoffs from each outcome.
- Additionally, in sequential games, the game stipulates the timing (or order) of moves.

COOPERATION

Game Theory

- How players should behave? ⇒ they should act rationally
- Which should be the ultimate result of a game?
  - Which is the player's *power*?
  - Which is the minimum payoff that a player can assure himself with his own resources?
  - Is it reasonable to think that other are hostile?

Game Theory

- Until which extent can agents do communicate?
- Players **may or may not** make agreements
- Can payoffs be shared among players? (that is , is it possible to pay third parties?)
- Which is the formal and causal relationship between actions and outcomes (payoff matrix)
- Which is the amount of information that agents have?

dmerida@lsi.upc.edu & ia@lsi.upc.edu

Defining Games

- Finite, n-person game <N, A, u>
  - N is a finite set of n players, indexed by i.
  - A = $A_1, \ldots, A_n$ is a set of actions for player i.
  - $u = \{u_1, \ldots, u_n\}$, a utility function for each player, where u_i: A➔ $\mathcal{R}$.

- Writing a 2-player game as a Matrix
  - Row player is player_1, column player is player_2.
  - Rows actions $a \in A_1$, and columns $a' \in A_2$.
  - Cells are outcomes, written as a tuple of utility values for each player

COOPERATION

Game Theory. An example

- A possible model for poker is the following:
  - Number of players= 2 ([X,Y])
  - Number of cards= 2 (*e.g.* [A, K])
  - Each player should put 1 euro in the bank

dmerida@lsi.upc.edu & ia@lsi.upc.edu

The name of the game

- Game Theory = Multi-person decision theory *or*
- Game Theory = Multi-agent decision theory

- The outcome of a game is determined by the actions *independently* taken by multiple decision makers.

- Strategic interaction.
  - Need to understand what the others will do
  - … what the others think that you will do.

Game Theory. An example

- A possible scenario for this game:
  - **Player X gets a card from the deck and after analyzing it either it can withdraw and lost the game (*i.e.* payoff = 0); or to bet 1 euro more.**
  - **Player Y can (without see X's card) withdraw (*i.e.* payoff = 0); or to bet 1 euro more.**
  - **Then X shows the card. If the card is A then X wins if not Y will win.**

  - Poker exemplifies a **zero-sum game** (ignoring the possibility of the house's cut), because one wins exactly the amount one's opponents lose.

COOPERATION

Game Theory. An example

	Y withdraw	Y bet
X withdraw	Payoff 0, 0 euro	Payoff -1, 1 euro
X bet	Payoff 1, -1 euro	A wins

A zero-sum game

dmerida@lsi.upc.edu & ia@lsi.upc.es

Khun Poker

- The deck includes only three playing cards, for example a King, Queen, and Ace.
-
- One card is dealt to each player, then the first player must bet or pass, then the second player may bet or pass. If any player chooses to bet the opposing player must bet as well (*call*) in order to stay in the round. After both players pass or bet the player with the highest card wins the **pot**.
- It is a zero sum two player game.

dmerida@lsi.upc.edu & ia@lsi.upc.edu

Prisoner's Dilemma

- Two partners in crime are separated into separate rooms at the police station and given a similar deal. If one implicates the other, he may go free while the other receives a life in prison. If neither implicates the other, both are given *moderate* sentences, and if both implicate the other, the sentences for both are *severe*.
- Each player has a *dominant strategy* to implicate the other, and thus in equilibrium each receives a harsh punishment, but both would be better off if each remained silent.
- In a repeated or iterated prisoner's dilemma, cooperation may be sustained through trigger strategies such as tit-for-tat.

Game Theory. Prisoner's Dilemma

	Y	X
X	a,a	b,c
Y	c,b	d,d

$c>a>d>b$

dmerida@lsi.upc.edu & ia@lsi.upc.es

Game Theory. Prisoner's Dilemma

	Y cooperate	X deny
X cooperate	Each serves 6 months	Prisoner X: 10 years Prisoner Y: goes free
Y deny	Prisoner X: goes free Prisoner Y: 10 years	Each serves 5 years

c>a>d>b

dmerida@lsi.upc.edu & ia@lsi.upc.es

Common payoff games

- A common payoff game is a game in which for all actions $a \in A_1 x \ldots x A_n$ and any pair of players i, j it is the case that $u_i(a) = u_j(a)$

- Common-payoff games are also called *pure coordination games or team games.* In such games the players have no conflicting (explicit) interests; their sole challenge is to *coordinate* on an action that is maximally beneficial to all.

dmerida@lsi.upc.edu & ia@lsi.upc.edu

Game Theory.

	Y	X
X	a,a	b,c
Y	c,b	d,d

a=d b=c

dmerida@lsi.upc.edu &
ia@lsi.upc.es

Game Theory.

	Y	X
X	1,1	0,0
Y	0,0	1,1

a=d b=c

dmerida@lsi.upc.edu &
ia@lsi.upc.es

COOPERATION

What is Game Theory?

- Study of rational behavior in ***interactive*** or ***interdependent*** situations

- **Bad news**:

  Knowing game theory does not guarantee winning

- **Good news**:

  Framework for thinking about strategic interaction

COOPERATION

Game Theory

- The game that nature seems to be playing is difficult to formulate. When different species compete, one knows how to define a **loss**: when a species dies out altogether, it loses, obviously. The defining **win**, however, is much more difficult because many coexist and will presumably for an infinite time; yet the humans in some sense consider themselves far ahead of the chicken, which will also allowed to go on to infinity.

  - S.Ulm

COOPERATION

Rules, strategies, payoffs and equilibrium

- Economic situations can be treated as games.

- The rules of the game determine **what**, **who** (can do) and **when** to do

- A player's **strategy** allows to create a plan for each situation .

dmerida@lsi.upc.edu & ia@lsi.upc.es

COOPERATION

Game Theory -Where?

Descriptive use

- Commerce: pricing …
- Political campaigns
- International conflicts
- Ecology: Biological equilibrium
- Negotiation: Tournaments
- Sociology: Study of mass behaviour
- Interpersonal conflicts: Divorces.
- Negotiation mechanisms
- IA: interactive computation, computational logic, implementation of rational behaviours

Normative use

dmerida@lsi.upc.edu & ia@lsi.upc.es

COOPERATION

Game Theory- Definitions

- N Players
- Utility function
- Possible actions (valid moves)
- Rules of interaction
- Prize/punishment rules*

- Each player aims to maximize its utility by choosing the right action (*do the right thing*) during the established interactions.

COOPERATION

Game Theory: Classification

- Representation
  - Normal Form (strategic form)
  - Extensive Form
- Types
  - Cooperative/non-Cooperative
  - Symmetric/ Asymmetric
  - Zero/non-Zero sum
  - Simultaneous/Sequential
  - Perfect Information/Imperfect information
  - Infinitely long games
  - Repeated/non-Repeated
  - 2 player/N player (N>2)

(PlayerA, PlayerB)	Opera	Football
Opera	(2,1)	(0,0)
Football	(0,0)	(1,2)

dmerida@lsi.upc.edu & ia@lsi.upc.es

Game Theory-Classification

- Representation
  - Normal Form (strategic form)
  - Extensive Form
- Types
  - Cooperative/non-Cooperative
  - Symmetric/Asymmetric
  - Zero/non-Zero sum
  - Simultaneous/Sequential
  - Perfect Information/Imperfect information
  - Incomplete/Complete information
  - Infinitely long games
  - Repeated/non-Repeated

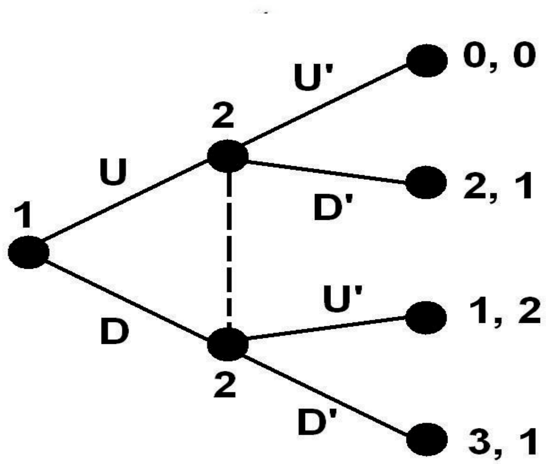

COOPERATION

dmerida@lsi.upc.edu & ia@lsi.upc.es

The Golden Rule

COMMANDMENT

Never assume that your opponents' behavior is fixed.

Predict their reaction to your behavior.

COOPERATION

MinMax

- In the zero-sum games the right combination of strategies allows to ***maximize*** own expectative and the ***minimization*** of our opponent.
- One can add a degree of pessimism into the application MinMax strategies add security.

dmerida@lsi.upc.edu & ia@lsi.upc.es

Mixed Strategy

- A strategy consisting of possible moves and a probability distribution (collection of weights) which corresponds to *how* frequently each move is to be played.

- A player would only use a ***mixed strategy*** when she is indifferent between several pure strategies, and when keeping the opponent guessing is desirable - that is, when the opponent can benefit from knowing the next move.

Min and Max

- $\max_i (\min_j M(i,j)) \leq \min_j (\max_i M(i,j))$
- Better to go second – one can react
- Proof
  - for any i'
  - $M(i',j') \leq \max_i M(i,j')$ for all j'
  - $\min_j M(i',j) \leq \min_j (\max_i M(i,j))$
  - $i' = \arg\max_i (\min_j M(i,j))$
  - $\max_i (\min_j M(i,j)) \leq \min_j (\max_i M(i,j))$

COOPERATION

Minimax Theorem

- *Every m-by-n 2-person zero-sum game has a solution. More precisely, there is a unique number **v**, called the value of the game, and there are optimal (pure or mixed) strategies **p*,q*** such that*

  $\max_p \min_q M(p,q) = \min_q \max_p M(p,q) = v = M(p^*,q^*)$

- i.e., *we know what's rational*

dmerida@lsi.upc.edu & ia@lsi.upc.es

Compromises

- Compromises are valuable.
- Compromises bring benefits, constraint (possible) actions and/or future options by modifying other agents' (possible) actions in our own benefit.

The benefit of compromises

- Agents may benefit of being able of limiting its future (possible) actions and perform those that they have compromised.
- An agent (only) gets compromises to perform a future action -- that constraints its future options/incentives – if it receives a larger gain.

Mechanisms to acquire mechanisms

- To compromise future actions/options is always a difficult decision.
- The law, social rules, *the rules of encountering*, the promises and honour rules do contribute to support agent to compromise.

Cooperation

- **Cooperation** is a type of coordination between agents that, in principle, are not opponents.
- The degree of success in cooperation is measured by the degree in which agents are capable in to maintain their own objectives allowing the others achieve their objectives.

Coordination

- An agent exists and performs its activity in a society in which other agents exist
-
- Coordination among agents is ***essential*** for achieving the ***goals*** and ***acting*** in a coherent manner.
- Coordination implies considering other agents' actions in the system when planning and executing one agent's actions.
- Coordination is also a means to achieve the coherent behaviour of the entire system

COOPERATION

Coordination

- **Coordination** is property of multiagent systems that ought to perform a task in a shared environment.

- The **degree of coordination** depends on:
  - The *necessity* of optimize resources
  - To avoid the paralysation of the process
  - To keep performance conditions

dmerida@lsi.upc.edu & ia@lsi.upc.es

Coordination

- An **activity** is the set of potential operations that an potential **actor,** that assumes a **role,** may perform to achieve a defined goal.
- An **actor** maybe an **agent** or a set of agents
- A set of **activities** and a given order of those is a **procedure**.

dmerida@lsi.upc.edu &
ia@lsi.upc.es

Coordination

- Coordination is a must in the implementation of MultiAgent Systems (MAS).

- Coordination becomes a critical element when the agents are heterogonous and autonomous.

COOPERATION

Coordination

- Coordination may imply ***cooperation*** and in this case the agent society works towards common goals to be achieved, but may also imply ***competition***, with agents having divergent or even antagonistic goals

dmerida@lsi.upc.edu & ia@lsi.upc.edu

Responsibility

- *People should be held responsible for the outcomes of exactly those choices that were free and unaffected by circumstances.* ***Aristotle***

Responsibility holders are decision-makers endowed with the capacity to foresee consequences of action (or inaction) and choose accordingly

Responsibility

- Responsibility implies ***deliberative*** capacity (free will, etc.): only autonomous and deliberative agents can be responsible for given events, both negative and positive.
  - It does not automatically imply (nor excludes) a decision to act or not act.

Negotiation

- *Negotiation* is as an iterative communication and decision making process between two or more agents who:
  - Cannot achieve their objectives through unilateral actions;
  - Exchange information comprising offers, counter-offers and arguments;
  - Deal with interdependent tasks; and
  - Search for a consensus which is a ***compromise*** decision
- There are two possible outcomes of a negotiation: a ***compromise*** or a ***disagreement***.

Coordination

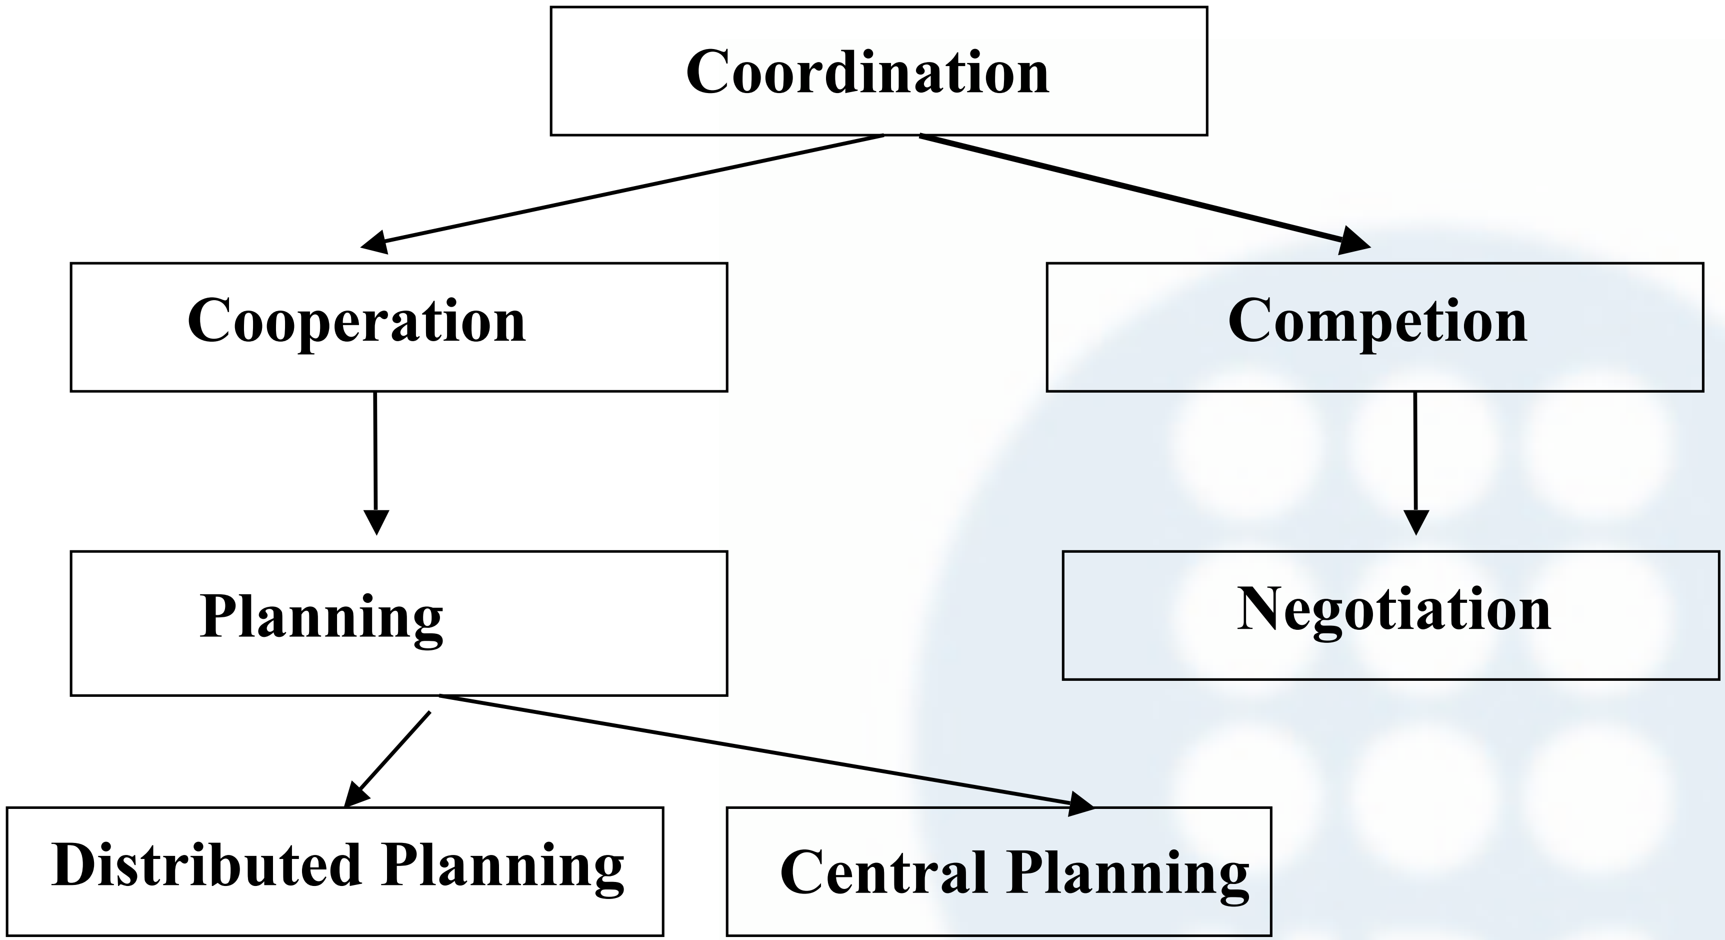

Teoría de Juegos

- The game that nature seems to be playing is difficult to formulate. When different species compete, one knows how to define a **loss**: when a species dies out altogether, it loses, obviously. The defining **win**, however, is much more difficult because many coexist and will presumably for an infinite time; yet the humans in some sense consider themselves far ahead of the chicken, which will also allowed to go on to infinity.

- S.Ulm

dmerida@lsi.upc.edu & ia@lsi.upc.es