

Strategic Games

Stéphane Airiau

Institute for Logic, Language & Computations (ILLC)
University of Amsterdam



12th European Agent Systems Summer School (EASSS 2010)
Ecole Nationale Supérieure des Mines de Saint-Etienne
Saint Etienne, France
August 23rd 2010.

- Agents have goals, they want to bring about some states of the world, they can take actions in their environment.
- In a multiagent system, agents interact, the actions of one may affect many other agents.
- How can we formally model such interactions?
- How should **rational** agents behave?
Game theory is one way.

Outline

- Today:** non-cooperative games
 - A central topic in Game theory: Strategic Games and Nash equilibrium.
 - Additional topics to provide a broader view of the field.
- Tomorrow:** cooperative games

Prisoner's dilemma

Two partners in crime, Row (R) and Column (C), are arrested by the police and are being interrogated in separate rooms. From Row's point of view, four different outcomes can occur:

- Only R confesses → R gets 1 year.
- Only C confesses → R spends 4 years in jail.
- Both confess → Both spend 3 years in prison.
- Neither one confesses → both get 2 years in prison.

The utility of an agent is (5 - number of years in prison).

	Column confesses	Column does not
Row confesses	2,2	4,1
Row does not	1,4	3,3

We can abstract this game and provide a generic game representation as follows:

Definition (Normal form game)

- A **normal form game (NFG)** is $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$ where
- N is the set of n players.
 - S_i is the set of strategies available to agent i .
 - $u_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}^n$ is the **payoff function** of agent i . It maps a **strategy profile** to a **utility**.

Terminology:

- an element $s = (s_1, \dots, s_n)$ of $S_1 \times \dots \times S_n$ is called a **strategy profile** or a **joint-strategy**.
- Let $s \in S_1 \times \dots \times S_n$ and $s'_i \in S_i$. We write (s'_i, s_{-i}) the joint-strategy which is the same as s except for agent i which plays strategy s'_i , i.e., $(s'_i, s_{-i}) = (s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$

What would you do?

- $N = \{Row, Column\}$
- $S_{Row} = S_{Column} = \{cooperate, defect\}$
- u_{Row} and u_{Column} are defined by the following bi-matrix.

Row \ Column	defect	cooperate
defect	2,2	4,1
cooperate	1,4	3,3

- Wait to know the other action?
- Not confess?
- Confess?
- Toss a coin?

Can you use some general principles to explain your choice?

Definition (strong dominance)

A strategy $x \in S_i$ for player i (**strongly**) **dominates** another strategy $y \in S_i$ if independently of the strategy played by the opponents, agent i (strictly) prefers x to y , i.e. $\forall s \in S_1 \times \dots \times S_n, u_i(x, s_{-i}) > u_i(y, s_{-i})$

Prisoner's dilemma

	C confesses	C does not
R confesses	2,2	4,1
R does not	1,4	3,3

Both players have a dominant strategy: to confess! From Row's point of view:

- if C confesses: R is better off confessing as well.
- if C does not: R can exploit and confess.

Battle of the sexes

	LS	RO
TO	2,2	4,3
BS	3,4	1,1

- Problem:** Where to go on a date: Soccer or Opera?
- Requirements:**
 - have a date!
 - be at your favourite place!

Do players have a dominant strategy?

Definition (Best response)

A strategy s_i of a player i is a **best response** to a joint-strategy s_{-i} of its opponents iff

$$\forall s'_i \in S_i, u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}).$$

Definition (Nash equilibrium)

A joint-strategy $s \in S_1 \times \dots \times S_n$ is a **Nash equilibrium** if each s_i is a best response to s_{-i} , that is

$$(\forall i \in N) (\forall s'_i \in S_i) u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

Battle of the sexes possesses two Nash equilibria (O, S) and (S, O) .

A **Nash equilibrium** is a joint-strategy in which no player could improve their payoff by unilaterally deviating from their assigned strategy.

Prisoner's dilemma

	C confesses	C does not
R confesses	2,2	2,1
R does not	1,4	3,3

Unique Nash equilibrium: both players confess!

- if R changes unilaterally, R loses!
- if C changes unilaterally, C loses!

Definition (Pareto optimal outcome)

A joint-strategy s is a **Pareto optimal outcome** if for no joint-strategy s'

$$\forall i \in N \ u_i(s') \geq u_i(s) \text{ and } \exists i \in N \ u_i(s') > u_i(s)$$

A joint-strategy is a Pareto optimal outcome when there is no outcome that is better for all players.

Prisoner's dilemma: Remaining silent is Pareto optimal.

discussion: It would be **rational** to confess! This seems counter-intuitive, as both players would be better off by keeping silent.

→ There is a conflict: the **stable** solution (i.e., the Nash equilibrium) is **not efficient**, as the outcome is not Pareto optimal.

Chicken game

In *Rebel Without a Cause*, James Dean's character's, Jim, is challenged to a "Chickie Run" with Buzz, racing stolen cars towards an abyss. The one who first jumps out of the car loses and is deemed a "chicken" (coward).

	Jim drives on	Jim turns
Buzz drives on	-10,-10	5,0
Buzz turns	0,5	1,1

Dominant Strategy? ✗

Nash equilibrium? ✗

Nash equilibrium: a summary

- When there is no dominant strategy, an equilibrium is the next best thing.
- A game may not have a Nash equilibrium.
- If a game possesses a Nash equilibrium, it may not be unique.
- Any combinations of dominant strategies is a Nash equilibrium.
- A Nash equilibrium may not be Pareto optimal.
- Two Nash equilibria may not have the same payoffs

Definition (Mixed strategy)

A mixed strategy p_i of a player i is a probability distribution over its strategy space S_i .

Assume that there are three strategies: $S_i = \{1,2,3\}$. Player i may decide to play strategy 1 with a probability of $\frac{1}{3}$, strategy 2 with a probability of $\frac{1}{2}$ and strategy 3 with a probability of $\frac{1}{6}$. The mixed strategy is then denoted as $\left\langle \frac{1}{3}, \frac{1}{2}, \frac{1}{6} \right\rangle$.

Given a mixed strategy profile $p = \langle p_1, \dots, p_n \rangle$, the expected utility for agent i is computed as follows:

$$E_i(p) = \sum_{s \in S_1 \times \dots \times S_n} \left(\prod_{j \in N} p_j(s_j) \right) \times u_i(s)$$

Battle of the sexes

	y	1-y	
	L	R	
x	T	2,2	4,3
1-x	B	3,4	1,1

The expected utility for the Row player is:
 $xy \cdot 2 + x(1-y) \cdot 4 + (1-x)y \cdot 3 + (1-x)(1-y) \cdot 1$
 $= -4xy + 3x + 2y + 1$

Given a mixed strategy profile $p = \langle p_1, \dots, p_n \rangle$, we write $\langle p'_i, p_{-i} \rangle$ the mixed strategy profile which is the same as p except for player i which plays mixed strategy p'_i , i.e., $\langle p'_i, p_{-i} \rangle = \langle p_1, \dots, p_{i-1}, p'_i, p_{i+1}, \dots, p_n \rangle$.

Definition (Mixed Nash equilibrium)

A **mixed Nash equilibrium** is a mixed strategy profile p such that $E_i(p) \geq E_i(p'_i, p_{-i})$ for every player i and every possible mixed strategy p'_i for i .

Battle of the sexes

	L	R
T	2,2	4,3
B	3,4	1,1

Let us consider that each player plays the mixed strategy $\left\langle \frac{3}{4}, \frac{1}{4} \right\rangle$.
None of the players have an incentive to deviate:

$$E_{row}(T) = \frac{3}{4} \cdot 2 + \frac{1}{4} \cdot 4 = \frac{5}{2} \quad E_{row}(B) = \frac{3}{4} \cdot 3 + \frac{1}{4} \cdot 1 = \frac{5}{2}$$

(players are indifferent)

Theorem (J. Nash, 1951)

Every finite strategic game has got at least one mixed Nash equilibrium.

note: The proofs are non-constructive and use Brouwer's or Kakutani's fixed point theorems.

J.F. Nash. Equilibrium points in n -person games. in *Proc. National Academy of Sciences of the United States of America*, 36:48-49, 1950.

Computing a Nash equilibrium

Complexity: In general, it is a hard problem. It is a PPA-complete problem.

Daskalakis, Goldberg, Papadimitriou: **The complexity of computing a Nash equilibrium**, in *Proc. 38th Ann. ACM Symp. Theory of Computing (STOC)*, 2006

There are complexity results and algorithms for different classes of games. We will not treat them in this tutorial.

Y. Shoham & K. Leyton-Brown: **Multiagent Systems**, Cambridge University Press, 2009. (Chapter 4)
Nisan, Roughgarden, Tardos & Vazirani: **Algorithmic Game Theory**, Cambridge University Press, 2007. (chapters 2, 3)

Other types of solution concepts for NFGs

Safety strategy

With Nash equilibrium, we assumed that the opponents were **rational agents**. What if the opponents are potentially **malicious**, i.e., their goal could be to minimize the payoff of the player?


Definition (Maxmin)

For player i ,

the **maxmin strategy** is $\operatorname{argmax}_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$,

and its **maxmin value** or **safety level** is $\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$.

- 1) player i chooses a (possibly mixed) strategy.
 - 2) the opponents $-i$ choose a (possibly mixed) strategy that *minimize* i 's payoff.
- the maxmin strategy *maximizes* i 's **worst case** payoff.



Whatever Column does, Row can guarantee itself a payoff of 2.5 by playing the mixed strategy $(\frac{1}{2}, \frac{1}{2})$.

Punish

Definition (Minmax)

For player i in a 2-player game,

the **minmax strategy** is $\operatorname{argmin}_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$,

and its **minmax value** is $\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$.

Player i 's strategy against player $-i$ in a 2-player game is a strategy that *minimizes* $-i$'s best-case payoff

Proposition

For a player i ,

$$\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \leq \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$$

Minimax theorem

Theorem

Minimax theorem (von Neumann, 1928)

In any finite two-player zero-sum game, for each player i , the maxmin strategy and minmax strategies are the same and are a Nash equilibrium of the game.

Minimax regret

Instead of assuming the opponents are rational (Nash equilibrium) or malicious (minimax), one can assume the **opponent is unpredictable** → avoid **costly mistakes**/minimize their worst-case losses.

	L	R
T	100,100	0,0
B	0,0	1,1

(T, L) is preferred by both agents.

However, (B, R) is also a NE.

There is no dominance.

How to explain that (T, L) should be preferred?

One can build a **regret-recording** game where the payoff function r_i is defined by $r_i(s_i, s_{-i}) = u_i(s_i^*, s_{-i}) - u_i(s_i, s_{-i})$, where s_i^* is i 's best response to s_{-i} , i.e., $r_i(s_i, s_{-i})$ is i 's **regret to have chosen s_i instead of s_i^*** .

$r_i \setminus r_j$	L	R
T	0,0	1,100
B	100,1	0,0

We define $\operatorname{regret}_i(s_i)$ as the maximal regret i can have from choosing s_i .

A **regret minimization strategy** is one that **minimizes the regret** function.

Correlated equilibrium

Battle of the sexes

	L	R
T	2,2	4,3
B	3,4	1,1

How to avoid the bad outcomes in which the agents fail to coordinate?

💡 **idea**: using a public random variable.

Example: the night before, the couple may condition their strategies based on weather (in the Netherlands, it is raining with a probability of 50%) as follows:

- if it rains at 5pm, we go to opera, otherwise, we go to football.
- both players increase their expected utility
- maybe a fairer solution

Definition (Correlated equilibrium)

Given an n -agent game $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, a **correlated equilibrium** is a tuple (ν, π, σ) , where

- ν is a tuple of random variables $\nu = (\nu_1, \dots, \nu_n)$ with respective domains $D = (D_1, \dots, D_n)$,
- π is a joint-distribution over ν ,
- $\sigma = (\sigma_1, \dots, \sigma_n)$ is a vector of mappings $\sigma_i : D_i \rightarrow S_i$,
- and for each agent i and every mapping $\sigma_i' : D_i \rightarrow S_i$ it is the case that

$$\sum_{d \in D} \pi(d) u_i(\sigma_1(d), \dots, \sigma_i(d), \dots, \sigma_n(d)) \geq \sum_{d \in D} \pi(d) u_i(\sigma_1(d), \dots, \sigma_i'(d), \dots, \sigma_n(d)).$$

Theorem

For every Nash equilibrium, there exists a corresponding correlated equilibrium.

Proof

Let s^* be a Nash equilibrium. We define

- $D_i = S_i$: strategy space and the domains of the random variables are the same.
- $\pi(d) = \prod_{i \in N} s_i^*(d_i)$
- $\sigma_i : D_i \rightarrow S_i, d_i \mapsto s_i^*$.


□

- Since a Nash equilibrium always exists, a correlated equilibrium always exists as well.
- However, a correlated equilibrium may not be a Nash equilibrium
- correlated equilibrium is a generalization of Nash equilibrium.

- We have considered games where each player choose their action **simultaneously**, and we have studied the normal-form representation.
- They are many games which rely on turn-taking, e.g., chess, card games, etc. Game theory has something to say about these games as well.
- We now introduce the **extended-form games (EFGs)**, in which a game is represented using a **tree** structure

Extended Form Games (EFGs)

Perfect-information game




A game is described by a **game tree**.

- the leaf nodes contain the payoff to the agents.
- the non-leaf nodes are **choice nodes**, labeled with the agent that make the decision for the node.
- The game tree is **common knowledge** before the agents start to play.
- During the play, the agents know which actions have been chosen in the past: this is called the **perfect information** case.

A **strategy** is a complete plan of actions of a player: a strategy specifies an action for each of its choice node.

ex: Player 1 decides for two nodes and has four strategies: (Left, Left), (Left, Right), (Right, Left) and (Right, Right).

Perfect-information game



Backward induction: when an agent knows the payoff at each of a node's children, it can decide the best action of the player making the decision for this node.

If there are ties, then how they are broken affects what happens higher up in the tree
= Multiple equilibria...

From an EFG to a NFG

- There can be an exponential number of pure strategies.
- Pure-strategy Nash equilibria of this game are (LL, LR), (LR, LR), (RL, LL), (RR, LL)
- But the only backward induction solution is (RL, LL)
- Nash equilibrium may be too weak for EFGs.


	L_1L_2	L_1R_1	R_1L_2	R_1R_2
L_1L_2	2,4	2,4	5,3	5,3
L_1R_1	2,4	2,4	5,3	5,3
R_1L_2	3,2	1,0	3,2	1,0
R_1R_2	3,2	0,1	3,2	0,1

Definition (Subgame)

A **subgame** is any sub-tree of the game tree.

Definition (Subgame-perfect equilibrium)

A strategy profile s is a **subgame-perfect equilibrium** for an EFG G iff for any subgame g , the restriction of s to g is a Nash equilibrium of g .



Other models of games

- Congestion games:** a special game which always possess a **pure** strategy Nash equilibrium
- Repeated games:** a NFG is played repeatedly (finitely/infinitely many times).
- Stochastic games:** uncertainty about the next game to play
- Bayesian games:** uncertainty about the current game

A **congestion game** is a tuple $(N, R, (S_i)_{i \in N}, (c_r)_{r \in R})$ where:

- $N = \{1, \dots, n\}$ is the set of **players**
- $R = \{1, \dots, m\}$ is the set of facilities or **resources**
- $S_i \subseteq M \setminus \emptyset$ denotes the set of **strategies** of player $i \in N$.
- $c_r(k)$ is the **cost** related to each user of resource $r \in M$ when exactly k players are using it.

Theorem

Every finite congestion game has a pure strategy Nash equilibrium.

R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria, in *International Journal of Game Theory*, 1973.

Theorem

Every congestion game is a potential game and every finite potential game is isomorphic to a congestion game

D. Monderer and L. S. Shapley **Potential Games**, in *Games and economic behavior*, 1996.

Repeated games

		Prisoner's dilemma	
		Defect	Cooperate
Cooperate	Defect	2,2	4,1
	Cooperate	1,4	3,3

When players are **rational**, both players confess!
If they trusted each other, they could both not confess and obtain (3,3).
If the same players have to repeatedly play the game, then it could be rational not to confess.

- One shot games:** there is no tomorrow. This is the type of games we have studied thus far.
- Repeated games:** model a likelihood of playing the game again with the same opponent. The NFG (N, S, u) being repeated is called the **stage game**.
 - finitely repeated games \rightarrow represent using a EFG and use backward induction to solve the game.
 - infinitely repeated games: the game tree would be infinite, use different techniques.

Infinitely repeated games

What is a strategy? In a repeated game, a **pure strategy** depends also on the **history** of play thus far.

ex: Tit-for-Tat strategy for the prisoner's dilemma:
Start by not confessing. Then, play the action played by the opponent during the previous iteration.

What is the players' objective?



- Average criterion:** Average payoff received throughout the game by player i : $\lim_{t \rightarrow \infty} \frac{\sum_{s=1}^t u_i(s^t)}{t}$, where s^t is the joint-strategy played during iteration t .
- Discounted-sum criterion:** Discounted sum of the payoff received throughout the game by player i : $\sum_{t=0}^{\infty} \gamma^t u_i(s^t)$, where γ is the discount factor (γ models how much the agent cares about the near term compared to long term).

Theorem (A Folk theorem)

Using the average criterion, any payoff vector v such that

- v is **feasible**, i.e., $\exists \lambda \in [0, 1]^{\prod_{i \in N} |S_i|}$ s.t. $v_i = \sum_{s \in \prod_{i \in N} S_i} \lambda_s v_i(s)$
- v is **enforceable** $v_i \geq \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$

can be sustained by a Nash equilibrium.



- In repeated games, the **same stage game** was played repeatedly.
- A **Stochastic game** is a set of NFGs. The agents **repeatedly** play games from this set. The next game is chosen with a probability which depends on the current game and the joint-action of the players.

Definition (Stochastic games)

A stochastic game is tuple $(N, \{S_i\}_{i \in N}, Q, P, \{u_i\}_{i \in N})$ where

- N is the set of players
- S_i is the strategy space of player i
- Q is a set of NFGs $q = (N, \{S_i\}_{i \in N}, \{v_i^q\}_{i \in N})$
- $P: Q \times \prod_{i \in N} S_i \times Q \rightarrow [0, 1]$ is the **transition function**. $P(q, s, q')$ is the probability that game q' is played after game q when the joint-strategy s was played in game q .
- $u_i: Q \times \prod_{i \in N} S_i$ is the **payoff function**. $u_i(q, s)$ is the payoff obtained by agent i when the joint-strategy s was played in game q .

In the definition, for ease of presentation, we assume that all the games have the same strategy space, which is not required.

- For stochastic games, the players know which game is currently played, i.e., they know the players of the game, the actions available to them, and their payoffs.

◦ In Bayesian games,

- there is **uncertainty** about the game currently being played.
- players have private information about the current game. The definition uses **information set**.

Definition (Bayesian game)

A **Bayesian game** is a tuple $(N, \{S_i\}_{i \in N}, G, P, \{I_i\}_{i \in N})$:

- N is the set of players.
- S_i is the set of strategies for agent i .
- G is a set of NFGs $g = (N, \{S_i\}_{i \in N}, \{u_i^g\}_{i \in N})$.
- P is a **common prior** over all games in G .
- I_i is the information set of agent i (a partition of G). A player knows the set which includes the current game, she does not know, however, which game it is in the set.
ex: G is composed of six games, $I_2 = \{\{g_1, g_3, g_4\}, \{g_2, g_5\}\}$. Agent 2 knows the current game is in $\{g_1, g_3, g_4\}$, but she does not know whether the game is g_1, g_3 , or g_4 .

Evolutionary game theory

- Models organisms in a large population (supposed infinite)
- **two** organisms are drawn randomly and play a 2-player game.
- the payoffs are linked to the fitness of the agents, and then, to their ability to reproduce.
- when an organism reproduces, a child adopt the same strategy as its parent.
- **Goal:** Are the strategies used by the organisms resilient to small mutant invasions? I.e. Is a strategy robust to evolutionary pressures?
→ evolutionary stability.

J. W. Weibull, Evolutionary game theory, the MIT press, 1997

Summary and Concluding remarks

When does an agent play?

- Agents play **simultaneously** (Rock/Paper/Scissors) → NFGs
- Agent play **sequentially** (chess, card games) → EFGs

What is known?

- **Complete information** games: the structure of the game and the preference of the agents are common knowledge.
- **Incomplete information games**
 - does a player know the preference of its opponents?
→ uncertainty, learning in games.
 - What kind of opponents? Rational? Malicious?
→ Nash equilibrium, minmax, maxmin, regret.

What can be observed? Are the agents able to observe the actions of the opponents (**perfect/imperfect information**)

How does the game develop?

- Is it a one stage game?
- Are there multiple stages? (repeated games) Does the structure of the game change? → Stochastic, Bayesian games
- Is the game played forever? → Infinitely repeated games

Nobel Laureates

1972	Arrow	Social choice
1994	Nash, Selten and Harsanyi	Game theory
1996	Vickrey	Mechanism design
1998	Sen	Social choice
2005	Schelling and Aumann	Game theory
2007	Hurwicz, Maskin and Myerson	Mechanism design

- **Game theory:** mathematical study of interaction among independent, self-interested agents. (Two sessions at AAMAS-10)
 - non-cooperative games
 - cooperative games
 - games with sequential actions
 - evolutionary game theory
- **Mechanism design:** study of protocol design for strategic agents (one session at AAMAS-09)
- **Social choice:** study of preference aggregation / collective decision making. (One session at AAMAS-10)

Resources

- Martin J. Osborne and Ariel Rubinstein. **A course in Game Theory**, the MIT Press, 1994. (freely available online)
- Yoav Shoham and Kevin Leyton-Brown. **Multiagent Systems**, Cambridge University Press, 2009
- Michael Wooldridge. **An Introduction to Multiagent Systems**, Wiley, 2009
- Noam Nisan, Tim Roughgarden, Éva Tardos & Vijay V. Vazirani. **Algorithmic Game Theory**, Cambridge University Press, 2007.
- gametheory.net

Tomorrow

Cooperative games

When agents work together, the group of agents, as a whole, gets a payoff.

- What groups of agents to form?
- How to distribute the payoff to the individual agents?