Game Theory 2025

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Game Theory

Game theory is the study of mathematical models to analyse strategic interactions between rational agents.

Example: Split or Steal

The *split-or-steal* game in the British television show "Golden Balls" is a good example for a game:

https://bit.ly/golden-balls-8200

Some of the main keywords we'll use in this course:

	Split	Steal
Split	$4k$, $4k$	0, $8k$
Steal	$8k$, 0	0, 0

- The *normal form* of this *strategic* (a.k.a. *noncooperative*) game is shown on the right.
- This is a one-shot game. Other games (like chess) can also be modelled using the *extensive form* (as a "game tree").
- The producers of the show engaged in *mechanism design:* refining the rules of the game to incentivise players to be "interesting".
- In a *coalitional* (a.k.a. *cooperative*) game, we might instead ask players to find a split that fairly reflects individual contributions.

Why?

Game theory plays a role in all of the academic disciplines that are covered by the *Master of Logic*. Examples:

- *Logic:* epistemic logics for modelling the reasoning patterns of agents engaging in strategic interaction
- *Philosophy:* systematic analysis of the conflicts arising between what people *ought* to do and what they *actually* do (ethics)
- *Linguistics:* signalling games as a model to explain linguistic conventions (game-theoretic pragmatics)
- *Mathematics:* infinite games (set theory)
- *Computer Science:* computational complexity of computing the equilibria of a game, to predict what the outcome might be

Why?

Game theory entered *AI* when it became clear that we can use it to study interaction between the software agents in a *multiagent system*. Nowadays, the study of "*economic paradigms*" is all over AI.

The influential *One Hundred Year Study on AI* (2016) singles out the following eleven "*hot topics*" in AI:

large-scale machine learning | deep learning | reinforcement learning | robotics | computer vision | natural language processing | collaborative systems | crowdsourcing and human computation | algorithmic *game theory* and computational social choice | internet of things | neuromorphic computing

P. Stone et al. "Artificial Intelligence and Life in 2030". One Hundred Year Study on Artificial Intelligence. Stanford, 2016.

Why?

Game theory also is one of the core tools of the trade for anyone conducting research in *computational social science*.

Course Organisation

Here is an overview of the *topics* to be covered in the course:

- Strategic games in normal form (3 weeks)
- Strategic games in extensive form (1 week)
- Mechanism design (1 week)
- Coalitional games (1 week)
- Guest lecture on current topics (final week)

To remain relevant to all of the diverse applications of game theory, the course will mostly focus on the mathematical properties of games. Thus, *mathematical maturity* (ability to handle proofs) is expected.

To better understand the *real-world relevance* of the game-theoretical concepts discussed we'll play some games in class ($\hookrightarrow$ *Game Academy*).

Read the *course manual* to find out about rules and practical matters.

Tutorials and Homework

During *tutorials* you are going to work on simple exercises designed to reinforce the material taught during lectures.

Most *homework* exercises will be of the problem-solving sort, requiring:

- a good *understanding* of the topic to see what the question is
- some *creativity* to find the solution
- *mathematical maturity*, to write it up correctly, often as a proof
- good *taste*, to write it up in a reader-friendly manner

<u>Also:</u> a small number of (optional) *programming assignments*.

Of course, solutions should be *correct*. But just as importantly, they should be *short* and *easy to understand*. (This is the advanced level: it's not anymore just about you getting it, it's now about your reader!)

Literature and Coverage

The course is largely based on Leyton-Brown and Shoham's *Essentials of Game Theory* (2008), which you'll need access to. But we'll *skip:*

- some of the "further solution concepts" in Chapter 3
- sequential equilibria (of imperfect-information games, in Chapter 5)
- repeated and stochastic games (all of Chapter 6)

On the other hand, we will go *beyond* the *Essentials* in other respects:

- material on congestion games, fictitious play, mechanism design
- more material on coalitional games (than what's in Chapter 8)
- proofs for most theorems

Of course, we cannot cover everything of interest. The most prominent omission might be *evolutionary game theory*.

K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan & Claypool Publishers, 2008.

Plan for Today

The remainder of today is an introduction to so-called strategic games in normal form. We are going to see:

- examples for and formal definition of *normal-form games*
- a definition of *stability* of an outcome (rational for all individuals)
- a definition of *efficiency* of an outcome (good for the group)

This (and more) is also covered in Chapters 1 and 2 of the *Essentials*.

We are also going to play a couple of games.

K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan & Claypool Publishers, 2008. Chapters 1 & 2.

The Prisoner's Dilemma

Two hardened criminals, Rowena and Colin, got caught by police and are being interrogated in separate cells. The police only has evidence for some of their minor crimes. Each is facing this dilemma:

- If we cooperate (C) and don't talk, then we each get 10 years for the minor crimes.
- If I cooperate but my partner defects (D) and talks, then I get 25 years.
- If my partner cooperates but I defect, then I go free (as crown witness).
- If we both defect, then we share the blame and get 20 years each.

	C	D
C	-10 / -10	-25 / 0
D	0 / -25	-20 / -20

What would you do? Why?

Let's Play: Prisoner's Dilemma Game

Here is the "same" game as before, but with simplified payoffs:

	C	D
C	₲15, ₲15	₲0, ₲25
D	₲25, ₲0	₲5, ₲5

We will try several variants:

- *pre-game communication* forbidden or allowed
- *one-shot* or *iterated* games, with (un)known number of iterations

For the iterated variant, your receive your average payoff (rounded).

<u>Soon:</u> *Specify a strategy (program) for how to play the iterated game.*

Real-World Relevance

Variants of the Prisoner's Dilemma (often with more than two players) commonly occur in real life. Examples:

- firms cooperating by not aggressively competing on price
- countries agreeing to caps on greenhouse gas emissions
- network users claiming only limited bandwidth

Strategic Games in Normal Form

A *normal-form game* is a tuple $\langle N, \boldsymbol{A}, \boldsymbol{u} \rangle$, where

- $N = \{1, \ldots, n\}$ is a finite set of *players* (or *agents*);
- $\boldsymbol{A} = A_1 \times \cdots \times A_n$ is a finite set of *action profiles* $\boldsymbol{a} = (a_1, \ldots, a_n)$, with A_i being the set of *actions* available to player i; and
- $\boldsymbol{u} = (u_1, \ldots, u_n)$ is a profile of *utility functions* $u_i : \boldsymbol{A} \to \mathbb{R}$.

Every player i chooses an action, say, a_i, giving rise to the profile $\boldsymbol{a}$. Actions are played *simultaneously*. Player i then receives payoff $u_i(\boldsymbol{a})$.

Remark: We use boldface italics to denote vectors (i.e., profiles) and Cartesian products (i.e., sets of profiles).

Nash Equilibria in Pure Strategies

Later we will allow players to randomise over actions. But today we restrict attention to *pure strategies:* strategy = action.

Notation: $(a'_i, \boldsymbol{a}_{-i})$ is short for $(a_1, \ldots, a_{i-1}, a'_i, a_{i+1}, \ldots, a_n)$.

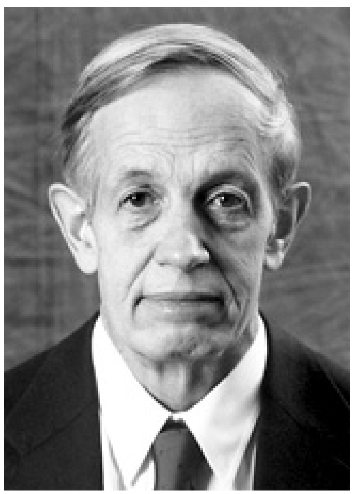

John F. Nash Jr.
(1928–2015)

We say that $a_i^\star \in A_i$ is a *best response* for player i to the (partial) action profile $\boldsymbol{a}_{-i}$, if $u_i(a_i^\star, \boldsymbol{a}_{-i}) \geqslant u_i(a'_i, \boldsymbol{a}_{-i})$ for all $a'_i \in A_i$.

We say that action profile $\boldsymbol{a} = (a_1, \ldots, a_n)$ is a *pure Nash equilibrium*, if a_i is a best response to $\boldsymbol{a}_{-i}$ for every player $i \in N$.

Thus, pure Nash equilibria are *stable* outcomes: no player has an incentive to unilaterally deviate from her assigned strategy.

Exercise: How Many Pure Nash Equilibria?

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

	L	R
T	2, 2	2, 2
B	2, 2	2, 2

	L	R
T	1, 2	2, 1
B	2, 1	1, 2

Pareto Efficiency

Next we formalise what we mean by "good for the group":

Vilfredo Pareto
(1848–1923)

Action profile $\boldsymbol{a}$ *Pareto-dominates* profile $\boldsymbol{a'}$, if $u_i(\boldsymbol{a}) \geqslant u_i(\boldsymbol{a'})$ for all players $i \in N$ and this inequality is strict in at least one case.

Action profile $\boldsymbol{a}$ is called *Pareto efficient*, if it is not Pareto-dominated by any other profile, i.e., if you cannot improve things for one player without harming any of the others.

Thus, the Prisoner's Dilemma illustrates a conflict between *stability* (both players defect) and *efficiency* (both players cooperate).

Let's Play: Numbers Game

Let's play the following game:

> *Every player submits a (rational) number between 0 and 100. We then compute the average (arithmetic mean) of all the numbers submitted and multiply that number with 2/3. Whoever got closest to this latter number wins the game.*

The winner gets ₲100. In case of a tie, the winners share the prize.

Summary

This has been a first introduction to game theory. We have seen:

- Definition of *normal-form games*
- *Nash equilibrium:* stable outcome for rational players
- *Pareto efficiency:* good (or rather: not bad) outcome for the group
- And: our idealised assumptions about players do not always match how people play in real life ($\hookrightarrow$ *behavioural game theory*)

Task: *Read the Course Manual (on Canvas). Ask next time if unclear.*

Task: *Compete in the Iterated Prisoner's Dilemma Tournament!*

What next? Mixed strategies, allowing players to randomise.