

Hannes Strass

Faculty of Computer Science, Institute of Artificial Intelligence, Computational Logic Group

Noncooperative Games in Normal Form

Lecture 1, 17th April 2023 // Algorithmic Game Theory, SS 2023

Introduction

A Story ...

- Two bank robbers are caught by the police and interrogated separately.
- The police tell each:
 - If you confess and incriminate your accomplice, then:
 - If your accomplice does not confess, then they will go to prison for 5 years; you will go free.
 - If your accomplice also confesses, then both of you will go to prison for 2 years.
 - Your accomplice gets the same offer.
- Each also knows that evidence is thin and if neither confesses, then both of them will go to prison for only 1 year.
- The bank robbers cannot coordinate their actions.
- What should they do?

And what does that have to do with games?

Terminology

Terminology

By **playing a game** we mean

- an **interaction** under preassigned rules,
- amongst one or more (typically several) **players**,
- each interested in **maximising** their gains,
- and acting **strategically** to this end.

A game is defined by its **rules**, which describe

- how the game is to be played,
- what each player is allowed to do (or not to do) in each situation,
- what single players can know about the current situation,
- when a game is over, and if so, who has won (and by how much).

Some Dimensions

Games can be distinguished along several dimensions:

single-player

vs. multi-player

perfect information

vs. imperfect information

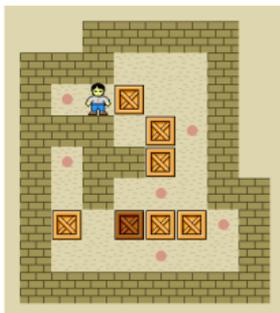
deterministic

vs. non-deterministic

sequential moves

vs. simultaneous moves

Examples: Sokoban, Chess, Poker, Rock-Paper-Scissors(-Lizard-Spock)



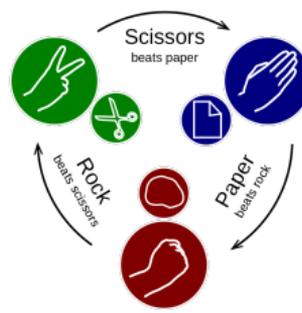
(C) Carloseow, CC BY 3.0



(C) Davidesouzari, CC BY-SA 4.0



(C) De Todd Klassy, CC BY 2.5



(C) Enzklop, CC BY-SA 3.0

Algorithmic Game Theory

- **Game Theory** studies mathematical models of strategic interactions between rational agents.
- Studied in mathematics, computer science, the social sciences, and economics.
- Fifteen game theorists have won the economics Nobel Prize.
- Historically: Started with two-player zero-sum games.
 - John von Neumann: *On the Theory of Games of Strategy* (1928)
 - John von Neumann and Oskar Morgenstern: *Theory of Games and Economic Behavior* (1944)
- **Algorithmic**: Special focus on algorithmic and computational aspects:
 - Algorithms and other approaches for obtaining or approximating solutions
 - Computational complexity of solution concepts

Utilities and Rationality

Terminology

The **utility** function of a player (also called **payoff** or **gain** function) models the player's preferences among different possible outcomes of a game.

Assumption

Players in a game act **rationaly**:

- Each player will choose moves that maximise their own utility function.
- Each player acts on their best estimates on other players' preferences.

↪ This does not necessarily imply that players are selfish; rather that possible altruistic motives have been built into the utility function.

↪ Being able to relate (sequences of) moves and utilities can be interpreted as players having access to “unlimited computational resources”.

Modelling and Fidelity

Ultimatum

There are two players, the **Proposer** and the **Responder**. The **Proposer** receives a certain amount of money, say 100. The **Proposer** can then offer any amount α of money to the **Responder**. If the **Responder** accepts the offer, the **Responder** gets the offered amount α and the **Proposer** keeps the remainder $100 - \alpha$. If the **Responder** rejects the offer, neither of them gets any money.

Naïve Analysis

A simple formalisation of this game models utility only by monetary gain. A game-theoretic analysis then yields that the **Proposer** should offer as little as possible, and the **Responder** should accept any offer.

Problem

This is not what humans do. In experiments, **Proposers** offer up to 50 (are being fair) and **Responders** reject offers below 25 (punish unfairness).

The Prisoner's "Game"

The situation of the two bank robbers can be modelled as a game:

- The two bank robbers are the two players.
- The game has only one round, where players move simultaneously.
- Each player has only two possible moves: **remain silent** or **confess**.
- The "gains" in the end are as follows:

	R_2 remain silent	R_2 confess
R_1 remain silent	R_1 gets 1 year, R_2 gets 1 year	R_1 gets 5 years, R_2 gets 0 years
R_1 confess	R_1 gets 0 years, R_2 gets 5 years	R_1 gets 2 years, R_2 gets 2 years

What is a "good" way to play this game?

The Prisoner's "Game"

- R_1 tries to approach the game rationally. They reason as follows:
- I cannot know (before making my choice) what R_2 does.
- But I do know there are only two possible options for R_2 :
 1. R_2 remains silent.
 - Then if I remain silent, I get 1 year;
 - if I confess, I get 0 years. \rightsquigarrow I confess.
 2. R_2 confesses.
 - Then if I remain silent, I get 5 years;
 - if I confess, I get 2 years. \rightsquigarrow I confess.
- So no matter what R_2 does, confess is always better for me.
- R_2 reasons likewise.
- So both confess and go to prison for 2 years.
- But had they both stayed silent, they would only go to prison for 1 year!

Overview

Introduction

Noncooperative Games in Normal Form

Solution Concepts

Dominant Strategies

Pareto Optimality

Pure Nash Equilibria

Noncooperative Games in Normal Form

Strategic Games in Normal Form

Definition

A (noncooperative) **game in normal form** is a tuple $G = (P, \mathbf{S}, \mathbf{u})$ where

- $P = \{1, 2, \dots, n\}$ is a set of players,
- $\mathbf{S} = (S_1, S_2, \dots, S_n)$ is a tuple of sets of (pure) **strategies** (or moves),
- $\mathbf{u} = (u_1, u_2, \dots, u_n)$ is a tuple of **utility functions** $u_i: \mathcal{S} \rightarrow \mathbb{R}$ (payoff, gain).

A (pure) **strategy profile** is a tuple $\mathbf{s} = (s_1, s_2, \dots, s_n) \in S_1 \times S_2 \times \dots \times S_n = \mathcal{S}$.

Example

For the prisoner's dilemma, there are two players $P = \{1, 2\}$ and

- $S_1 = S_2 = \{\text{Silent}, \text{Confess}\},$
- $\mathcal{S} = \{(\text{Silent}, \text{Silent}), (\text{Silent}, \text{Confess}), (\text{Confess}, \text{Silent}), (\text{Confess}, \text{Confess})\},$
- $u_1 = \{(\text{Silent}, \text{Silent}) \mapsto 4, (\text{Silent}, \text{Confess}) \mapsto 0, (\text{Confess}, \text{Silent}) \mapsto 5, (\text{Confess}, \text{Confess}) \mapsto 3\},$
- $u_2 = \{(\text{Silent}, \text{Silent}) \mapsto 4, (\text{Silent}, \text{Confess}) \mapsto 5, (\text{Confess}, \text{Silent}) \mapsto 0, (\text{Confess}, \text{Confess}) \mapsto 3\}.$

Thus the utility is measured by the number of years (out of the next five) not spent in prison.

Representation of Two-Player Games

Typically, a two-player game $G = (\{1, 2\}, (S_1, S_2), (u_1, u_2))$ is represented by specifying the $|S_1| \times |S_2|$ -payoff matrix that lists the gains of each player for each possible strategy profile ($S_i = \{s_{i,1}, \dots, s_{i,k_i}\}$ for $i \in \{1, 2\}$):

(1, 2)	$s_{2,1}$	$s_{2,2}$...	s_{2,k_2}
$s_{1,1}$	$(u_1(s_{1,1}, s_{2,1}), u_2(s_{1,1}, s_{2,1}))$	$(u_1(s_{1,1}, s_{2,2}), u_2(s_{1,1}, s_{2,2}))$		$(u_1(s_{1,1}, s_{2,k_2}), u_2(s_{1,1}, s_{2,k_2}))$
$s_{1,2}$	$(u_1(s_{1,2}, s_{2,1}), u_2(s_{1,2}, s_{2,1}))$	$(u_1(s_{1,2}, s_{2,2}), u_2(s_{1,2}, s_{2,2}))$		$(u_1(s_{1,2}, s_{2,k_2}), u_2(s_{1,2}, s_{2,k_2}))$
\vdots			\ddots	
s_{1,k_1}	$(u_1(s_{1,k_1}, s_{2,1}), u_2(s_{1,k_1}, s_{2,1}))$	$(u_1(s_{1,k_1}, s_{2,2}), u_2(s_{1,k_1}, s_{2,2}))$		$(u_1(s_{1,k_1}, s_{2,k_2}), u_2(s_{1,k_1}, s_{2,k_2}))$

Example

The prisoner's dilemma is thus succinctly represented as

(1, 2)	Silent	Confess
Silent	(4,4)	(0,5)
Confess	(5,0)	(3,3)

Further Strategic Games (1)

Battle of the Partners

Two partners, **Cat** and **Dee**, discuss how to spend the evening. Each has their personal preference what to do, but overall they want to spend the evening together.

(Cat, Dee)	Concert	Cinema
Concert	(10,7)	(2,2)
Cinema	(0,0)	(7,10)

Chicken

Two people, **Eli** and **Fyn**, are racing towards each other in cars. Whoever swerves ("chickens out") loses face. If neither swerves, both get seriously injured.

(Eli, Fyn)	Swerve	RaceOn
Swerve	(2,2)	(1,3)
RaceOn	(3,1)	(0,0)

Further Strategic Games (2)

Penalties

Two football players face off at a (simplified) single penalty kick. The kicker can kick left or right; the goal keeper can jump left or right. The kicker scores a goal iff they choose a different side than the keeper.

(Kicker, Keeper)	JumpL	JumpR
KickL	(-1,1)	(1,-1)
KickR	(1,-1)	(-1,1)

Rock-Paper-Scissors

Each player chooses one of three symbols, each of which wins/loses against exactly one other symbol.

(Ann, Bob)	Rock	Paper	Scissors
Rock	(0,0)	(-1,1)	(1,-1)
Paper	(1,-1)	(0,0)	(-1,1)
Scissors	(-1,1)	(1,-1)	(0,0)

Solution Concepts

Dominant Strategies

Definition

Let $G = (P, \mathbf{S}, \mathbf{u})$ be a game in normal form.

A strategy $s_i \in S_i$ of player i is **dominant** (or **weakly dominant**) iff

$$u_i(s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_n) \geq u_i(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$$

for all strategies $s'_i \in S_i$ and all strategies $s_j \in S_j$ with $1 \leq j \leq n$ and $i \neq j$.

s_i is **strictly dominant** iff the inequality is strict for all $s'_i \neq s_i$ and for all s_j .

Example

- Strategy **Confess** is strictly dominant for every player of the prisoner's dilemma.
- There are no dominant strategies in the battle of the partners, chicken, penalties, and rock-paper-scissors.

Pareto Dominance and Pareto Optimality

Definition

Let $G = (P, \mathcal{S}, \mathbf{u})$ be a game in normal form and $\mathbf{s}, \mathbf{t} \in \mathcal{S}$ be strategy profiles.

1. **\mathbf{s} weakly Pareto-dominates \mathbf{t}** iff for all $1 \leq i \leq n$, we have $u_i(\mathbf{s}) \geq u_i(\mathbf{t})$.
2. **\mathbf{s} Pareto-dominates \mathbf{t}** iff \mathbf{s} weakly Pareto-dominates \mathbf{t} and there exists one $1 \leq j \leq n$ such that $u_j(\mathbf{s}) > u_j(\mathbf{t})$.
3. **\mathbf{s} strongly Pareto-dominates \mathbf{t}** iff for all $1 \leq i \leq n$, we have $u_i(\mathbf{s}) > u_i(\mathbf{t})$.
4. **\mathbf{t} is Pareto-optimal** iff there is no $\mathbf{s} \in \mathcal{S}$ that Pareto-dominates \mathbf{t} .
5. **\mathbf{t} is weakly Pareto-optimal** iff there is no $\mathbf{s} \in \mathcal{S}$ that strongly Pareto-dominates \mathbf{t} .

Intuitively: In a Pareto optimum, no player can unilaterally gain by switching strategies without some other player being worse off.

Pareto Optimality: Examples

Examples

- The strategy profile (**Silent**, **Silent**) is a Pareto optimum in the prisoner's dilemma; so are the profiles (**Silent**, **Confess**) and (**Confess**, **Silent**). The profile (**Silent**, **Silent**) strongly Pareto-dominates (**Confess**, **Confess**).
- Battle of the partners has two Pareto optima: (**Concert**, **Concert**) and (**Cinema**, **Cinema**).
- Similarly, chicken has the two Pareto optima (**Swerve**, **RaceOn**) and (**RaceOn**, **Swerve**), and a third Pareto optimum (**Swerve**, **Swerve**).
- All strategy profiles of penalties are Pareto-optimal.
- Similarly, all strategy profiles of rock-paper-scissors are Pareto-optimal.

Best Responses

Definition

Let $(G, \mathbf{S}, \mathbf{u})$ be a game in normal form and for a player $i \in P$, denote

- $\mathbf{S}_{-i} := S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n$, and
- $\mathbf{s}_{-i} := (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ for all $\mathbf{s} \in \mathcal{S}$.

1. A strategy $s_i \in S_i$ is **player i 's best response to \mathbf{s}_{-i}** iff for all strategies $s'_i \in S_i$: $\mathbf{u}_i(s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_n) \geq \mathbf{u}_i(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$.
2. s_i is the **strictly best response** to \mathbf{s}_{-i} iff s_i is the only best response to \mathbf{s}_{-i} .

- \mathbf{s}_{-i} is the strategy profile \mathbf{s} with out the strategy of player i .
- \mathbf{S}_{-i} is the set of all such strategy profiles.
- The best response to \mathbf{s}_{-i} is player i 's best move given the others play \mathbf{s}_{-i} .

Note

A dominant strategy is always a best response, but not vice versa.

Best Response: Examples

Examples

- In the prisoner's dilemma, for both players, **Confess** is a best response for **Silent** and **Confess** is also a best response for **Confess**.
- In the battle of the partners, the best response to **Concert** is **Concert**, and the best response to **Cinema** is **Cinema**.
- In chicken, the best response to **Swerve** is **RaceOn**, and the best response to **RaceOn** is **Swerve**.
- In penalties, the best response to **KickR** is **JumpR**, the best response to **JumpR** is **KickL**, the best response to **KickL** is **JumpL**, and the best response to **JumpL** is **KickR**.
- Similarly, in rock-paper-scissors, the best response to **Rock** is **Paper**, the best response to **Paper** is **Scissors**, and the best response to **Scissors** is **Rock**.

Nash Equilibrium (in Pure Strategies)

Definition

Let $(G, \mathbf{S}, \mathbf{u})$ be a game in normal form and let $i \in P$ be a player.

1. A strategy profile $\mathbf{s} \in \mathcal{S}$ is (in) a **Nash equilibrium in pure strategies** iff for all $1 \leq i \leq n$, strategy s_i is a best response for \mathbf{s}_{-i} .
2. \mathbf{s} is (in) a **strict Nash equilibrium in pure strategies** iff \mathbf{s} is the only strategy profile in a Nash equilibrium in pure strategies.

Pure Nash Equilibria: Examples

- The prisoner's dilemma has the single Nash equilibrium in pure strategies (**Confess, Confess**), where every player plays their dominant strategy.
- Battle of the partners has two Nash equilibria in pure strategies: (**Concert, Concert**) and (**Cinema, Cinema**); incidentally both Pareto optimal.
- Penalties does not have a Nash equilibrium in pure strategies.

Solution Concepts: Stocktaking

Theorem

Let $G = (P, \mathbf{S}, \mathbf{u})$ be a game in normal form. If every player $i \in P$ has a (strictly) dominant strategy $s_i \in S_i$, then $\mathbf{s} := (s_1, \dots, s_n)$ is a (strict) Nash equilibrium in pure strategies.

- Nash equilibrium is the “standard” solution concept for strategic games; it can be thought of as a prediction about how the game will be played.
- Dominant strategies are helpful, but rarely exist.
- Pareto optimality is useful as a normative goal rather than as a prediction.

Proposition

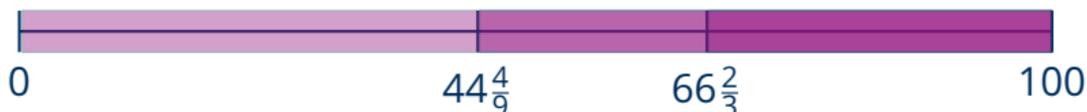
Let G be a finite noncooperative game in normal form where the payoffs are explicitly specified. The pure Nash equilibria of G can be computed in deterministic polynomial time by exhaustive search for best responses.

Limitations of Predictive Power

Guessing Numbers

A number n of players guess a real number $s_i \in [0, 100]$ each.
Whoever's guess comes closest to two thirds of the average guess wins.

Let (s_1, \dots, s_n) be a strategy profile and denote $a := \frac{1}{n} \cdot \sum_{i=1}^n s_i$ and $z := \frac{2}{3} \cdot a$.



- The maximally possible value for z is $z_{\max} = 66\frac{2}{3}$, dominating any $s > z_{\max}$.
- Rational players will not play dominated strategies (everyone knows ...).
- So the maximally possible value for z is actually $z'_{\max} = 44\frac{4}{9}, \dots$
- The game has $(0, \dots, 0)$ as strict Nash equilibrium in pure strategies.
- Experiments show that human players do not play the Nash equilibrium.
- Their guesses typically peak at 22, 33, and 0, less visibly at 67 and 100.

Conclusion

Summary

- Games can model real-life situations, but model fidelity is important.
- Noncooperative (strategic) games in **normal form** comprise players, **strategies** for the players, and **gain functions** for all **strategy profiles**.
- Various concepts can help predict/analyse the outcome of a game:
 - **Dominant strategies**
 - **Pareto optimality**
 - (pure) **Nash equilibria**
- We have analysed a number of example games: **prisoner's dilemma**, **battle of the partners**, **chicken**, **penalties**, and **guessing numbers**.
- Pure Nash equilibria need not always exist.

Action Points

- Find new interpretations (stories) for the example games.