Algorithmic game theory

Martin Balko

1st lecture

September 30th 2025

Basic info

Basic info

- Webpage: *https://kam.mff.cuni.cz/~balko/ath2526/ATH.html*
  - lecture info, topics covered, presentations, lecture notes ...

Basic info

- Webpage: *https://kam.mff.cuni.cz/~balko/ath2526/ATH.html*
    - lecture info, topics covered, presentations, lecture notes ...
- Recommended literature:
    - M. Balko: Algorithmic game theory: lecture notes.
    - The notes are still under construction. Comments are welcome.

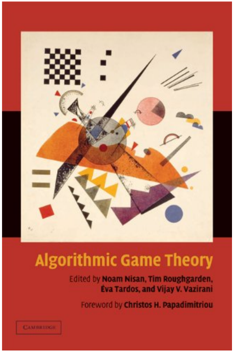

Figure: Algorithmic game theory by Nisan et al.

Source: https://amazon.com

Game theory

Game theory

- study of mathematical models of strategic interaction among rational decision-makers.

Zdroj: https://quantamagazine.org

Game theory

- study of mathematical models of strategic interaction among rational decision-makers.

Zdroj: https://quantamagazine.org

- We focus on the algorithmic side of the game theory.

Game theory

- study of mathematical models of strategic interaction among rational decision-makers.

Zdroj: https://quantamagazine.org

- We focus on the algorithmic side of the game theory.
- Several real-word applications.

Game theory

- study of mathematical models of strategic interaction among rational decision-makers.

Zdroj: https://quantamagazine.org

- We focus on the algorithmic side of the game theory.
- Several real-word applications.
- More than ten game theorists have won the Nobel Prize in economics.

Other game theory classes

Other game theory classes

- There are two other classes taught by people from our group.

Other game theory classes

- There are two other classes taught by people from our group.
- Modern Algorithmic Game Theory (NOPT021)

Other game theory classes

- There are two other classes taught by people from our group.
- Modern Algorithmic Game Theory (NOPT021)
  - Taught by Martin Schmid and Radovan Haluška,

Other game theory classes

- There are two other classes taught by people from our group.
- Modern Algorithmic Game Theory (NOPT021)
  - Taught by Martin Schmid and Radovan Haluška,
  - runs simultaneously with this class in the winter term; still a chance to enroll,

Other game theory classes

- There are two other classes taught by people from our group.
- Modern Algorithmic Game Theory (NOPT021)
  - Taught by Martin Schmid and Radovan Haluška,
  - runs simultaneously with this class in the winter term; still a chance to enroll,
  - covers similar topics as this class, plus some more,

Other game theory classes

- There are two other classes taught by people from our group.
- Modern Algorithmic Game Theory (NOPT021)
  - Taught by Martin Schmid and Radovan Haluška,
  - runs simultaneously with this class in the winter term; still a chance to enroll,
  - covers similar topics as this class, plus some more,
  - focuses on the implementation of the algorithms such as support enumeration or regret minimization.

Other game theory classes

- There are two other classes taught by people from our group.
- Modern Algorithmic Game Theory (NOPT021)
  - Taught by Martin Schmid and Radovan Haluška,
  - runs simultaneously with this class in the winter term; still a chance to enroll,
  - covers similar topics as this class, plus some more,
  - focuses on the implementation of the algorithms such as support enumeration or regret minimization.
- Advanced Modern Algorithmic Game Theory (NOPT022) (long title)

Other game theory classes

- There are two other classes taught by people from our group.

- Modern Algorithmic Game Theory (NOPT021)
  - Taught by Martin Schmid and Radovan Haluška,
  - runs simultaneously with this class in the winter term; still a chance to enroll,
  - covers similar topics as this class, plus some more,
  - focuses on the implementation of the algorithms such as support enumeration or regret minimization.

- Advanced Modern Algorithmic Game Theory (NOPT022) (long title)
  - A brand new course taught by Martin Schmid and Radovan Haluška,

Other game theory classes

- There are two other classes taught by people from our group.
- Modern Algorithmic Game Theory (NOPT021)
    - Taught by Martin Schmid and Radovan Haluška,
    - runs simultaneously with this class in the winter term; still a chance to enroll,
    - covers similar topics as this class, plus some more,
    - focuses on the implementation of the algorithms such as support enumeration or regret minimization.
- Advanced Modern Algorithmic Game Theory (NOPT022) (long title)
    - A brand new course taught by Martin Schmid and Radovan Haluška,
    - a continuation of the previous class that focuses on approximation methods and large games.

Other game theory classes

Other game theory classes

Novinky.cz

Novinky.cz » Internet a PC » Češi vytvořili umělou inteligenci, která drtí v … Podrubriky: Hardware • Software • Testy • Hry a herní systémy • Mobil

Češi vytvořili umělou inteligenci, která drtí v pokeru jednoho hráče za druhým

3. 3. 2017, 15:44 - mif, Novinky

Facebook Twitter

Vědci z Matematicko-fyzikální fakulty Univerzity Karlovy a Fakulty elektrotechnické ČVUT v Praze pracovali několik posledních měsíců na vývoji umělé inteligence, jejímž hlavním úkolem bude stát se špičkou v karetní hře Poker Texas Hold'em. A to se skutečně podařilo, program porazil hned několik profesionálních hráčů.

Further info and sylabus

Further info and sylabus

- You can earn up to 100 points for completing small weekly quizzes and homeworks.

Further info and sylabus

- You can earn up to 100 points for completing small weekly quizzes and homeworks.
- To receive credit, you need at least 65 points. If you earn at least 85 points, you will not be required to solve a problem during the exam.

Further info and sylabus

- You can earn up to 100 points for completing small weekly quizzes and homeworks.
- To receive credit, you need at least 65 points. If you earn at least 85 points, you will not be required to solve a problem during the exam.
- Preliminary plan:

Further info and sylabus

- You can earn up to 100 points for completing small weekly quizzes and homeworks.
- To receive credit, you need at least 65 points. If you earn at least 85 points, you will not be required to solve a problem during the exam.
- Preliminary plan:
  - Finding Nash equilibria
    - Nash equilibria and Nash's Theorem,
    - zero-sum games,
    - bimatrix games and the Lemke–Howson algorithm,
    - other notions of equilibria,
    - regret minimization.
  - Mechanism design,
    - auctions (Vickrey),
    - Myerson's lemma and its applications,
    - revenue maximization.

Finding Nash equilibria

Normal-form games

Normal-form games

- We use the following most fundamental representation of games.

Normal-form games

- We use the following most fundamental representation of games.
- A normal-form game is a triple (P, A, u), where

Normal-form games

- We use the following most fundamental representation of games.
- A normal-form game is a triple (P, A, u), where
  - P is a finite set of n players,

Normal-form games

- We use the following most fundamental representation of games.
- A normal-form game is a triple (P, A, u), where
  - P is a finite set of n players,
  - $A = A_1 \times \cdots \times A_n$ is a set of action profiles, where A_i is a set of actions available to player i,

Normal-form games

- We use the following most fundamental representation of games.
- A normal-form game is a triple (P, A, u), where
  - P is a finite set of n players,
  - $A = A_1 \times \cdots \times A_n$ is a set of action profiles, where A_i is a set of actions available to player i,
  - and $u = (u_1, \ldots, u_n)$ is an n-tuple, where each $u_i : A \to \mathbb{R}$ is the utility function for player i.

Normal-form games

- We use the following most fundamental representation of games.
- A normal-form game is a triple (P, A, u), where
  - P is a finite set of n players,
  - $A = A_1 \times \cdots \times A_n$ is a set of action profiles, where A_i is a set of actions available to player i,
  - and $u = (u_1, \ldots, u_n)$ is an n-tuple, where each $u_i \colon A \to \mathbb{R}$ is the utility function for player i.
- Knowing the utility function, all players i simultaneously choose an action a_i from A_i. The resulting action profile $a = (a_1, \ldots, a_n)$ is then evaluated using the utility function.

Normal-form games

- We use the following most fundamental representation of games.
- A normal-form game is a triple (P, A, u), where
  - P is a finite set of n players,
  - $A = A_1 \times \cdots \times A_n$ is a set of action profiles, where A_i is a set of actions available to player i,
  - and $u = (u_1, \ldots, u_n)$ is an n-tuple, where each $u_i \colon A \to \mathbb{R}$ is the utility function for player i.

- Knowing the utility function, all players i simultaneously choose an action a_i from A_i. The resulting action profile $a = (a_1, \ldots, a_n)$ is then evaluated using the utility function.
- The ith coordinate $u_i(a)$ of $u(a)$ is the gain of player i on a.

Normal-form games: Rock-Paper-Scissors

Normal-form games: Rock-Paper-Scissors

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

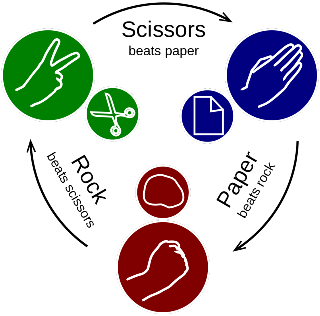

Sources: https://en.wikipedia.org/

Normal-form games: Rock-Paper-Scissors-Lizard-Spock

Normal-form games: Rock-Paper-Scissors-Lizard-Spock

	Rock	Paper	Scissors	Lizard	Spock
Rock	(0,0)	(-1,1)	(1,-1)	(1,-1)	(-1,1)
Paper	(1,-1)	(0,0)	(-1,1)	(-1,1)	(1,-1)
Scissors	(-1,1)	(1,-1)	(0,0)	(1,-1)	(-1,1)
Lizard	(-1,1)	(1,-1)	(-1,1)	(0,0)	(1,-1)
Spock	(1,-1)	(-1,1)	(1,-1)	(-1,1)	(0,0)

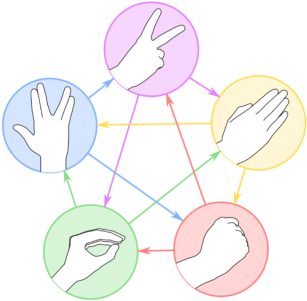

Sources: https://bigbangtheory.fandom.com/

- "Scissors cuts Paper, Paper covers Rock, Rock crushes Lizard, Lizard poisons Spock, Spock smashes Scissors, Scissors decapitates Lizard, Lizard eats Paper, Paper disproves Spock, Spock vaporizes Rock (and as it always has) Rock crushes Scissors."

Normal-form games: Rock-Paper-Scissors-Lizard-Spock

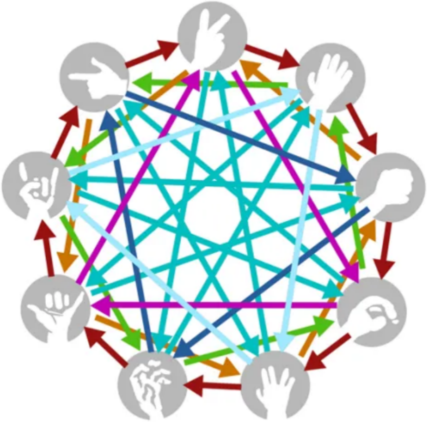

ROCK PAPER SCISSORS
LIZARD SPOCK
SPIDER-MAN BATMAN
WIZARD GLOCK

Scissors cuts paper.
Paper covers rock.
Rock crushes lizard.
Lizard poisons Spock.
Spock zaps wizard.
Wizard stuns Batman.
Batman scares Spider-Man.
Spider-Man disarms glock.
Glock breaks rock.
Rock interrupts wizard.
Wizard burns paper.
Paper disproves Spock.
Spock befuddles Spider-Man.
Spider-Man defeats lizard.
Lizard confuses Batman
(because he looks like Killer Croc).
Batman dismantles scissors.
Scissors cut wizard.
Wizard transforms lizard.
Lizard eats paper.
Paper jams glock.
Glock kills Batman's mom.
Batman explodes rock.
Rock crushes scissors.
Scissors decapitates lizard.
Lizard is too small for glock.
Glock shoots Spock.
Spock vaporizes rock.
Rock knocks out Spider-Man.
Spider-Man rips paper.
Paper delays Batman.
Batman hangs Spock.
Spock smashes scissors.
Scissors cut Spider-Man.
Spider-Man annoys wizard.
Wizard melts glock.
Glock dents scissors.

ROCK PAPER SCISSORS SPOCK LIZARD by Sam Kass and Karen Bryla, and then, Brian Yan messed it up into this.

Source: https://www.naturphilosophie.co.uk/

Normal-form games: Chess

Normal-form games: Chess

Source: https://edition.cnn.com/

- Chess as a normal-form game: Each action of player $i \in \{\text{black}, \text{white}\}$ is a list of all possible situations that can happen on the board together with the move player i would make in that situation. Then we can simulate the whole game of chess in one round.

Strategies

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
  - "select a single action and play it",

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
    - "select a single action and play it",
    - a pure-strategy profile is an n-tuple $(s_1, \dots, s_n)$, where $s_i \in A_i$ for each player i.

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
  - "select a single action and play it",
  - a pure-strategy profile is an n-tuple $(s_1, \ldots, s_n)$, where $s_i \in A_i$ for each player i.
- A mixed strategy s_i of player i is a probability distribution over A_i.

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
  - "select a single action and play it",
  - a pure-strategy profile is an n-tuple $(s_1, \ldots, s_n)$, where $s_i \in A_i$ for each player i.
- A mixed strategy s_i of player i is a probability distribution over A_i.
  - that is, s_i assigns a value $s_i(a_i) \in [0, 1]$ to each $a_i \in A_i$ so that

$$\sum_{a_i \in A_i} s_i(a_i) = 1.$$

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
  - "select a single action and play it",
  - a pure-strategy profile is an n-tuple $(s_1, \dots, s_n)$, where $s_i \in A_i$ for each player i.
- A mixed strategy s_i of player i is a probability distribution over A_i.
  - that is, s_i assigns a value $s_i(a_i) \in [0,1]$ to each $a_i \in A_i$ so that

$$\sum_{a_i \in A_i} s_i(a_i) = 1.$$

  - $s_i(a_i) =$ "probability that i chooses a_i as his action".

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
  - "select a single action and play it",
  - a pure-strategy profile is an n-tuple $(s_1, \ldots, s_n)$, where $s_i \in A_i$ for each player i.
- A mixed strategy s_i of player i is a probability distribution over A_i.
  - that is, s_i assigns a value $s_i(a_i) \in [0, 1]$ to each $a_i \in A_i$ so that

$$\sum_{a_i \in A_i} s_i(a_i) = 1.$$

  - $s_i(a_i) =$ "probability that i chooses a_i as his action".
  - We let S_i be the set of all mixed strategies of player i.

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
  - "select a single action and play it",
  - a pure-strategy profile is an n-tuple $(s_1, \ldots, s_n)$, where $s_i \in A_i$ for each player i.
- A mixed strategy s_i of player i is a probability distribution over A_i.
  - that is, s_i assigns a value $s_i(a_i) \in [0, 1]$ to each $a_i \in A_i$ so that

$$\sum_{a_i \in A_i} s_i(a_i) = 1.$$

  - $s_i(a_i) =$ "probability that i chooses a_i as his action".
  - We let S_i be the set of all mixed strategies of player i.
  - a mixed-strategy profile is an n-tuple $(s_1, \ldots, s_n)$, where $s_i \in S_i$ for each player i.

Strategies

- Each player i follows a certain strategy (a prescription how he chooses his actions from A_i).
- A pure strategy s_i of player i is an action from A_i.
  - "select a single action and play it",
  - a pure-strategy profile is an n-tuple $(s_1, \ldots, s_n)$, where $s_i \in A_i$ for each player i.
- A mixed strategy s_i of player i is a probability distribution over A_i.
  - that is, s_i assigns a value $s_i(a_i) \in [0, 1]$ to each $a_i \in A_i$ so that

$$\sum_{a_i \in A_i} s_i(a_i) = 1.$$

  - $s_i(a_i) =$ "probability that i chooses a_i as his action".
  - We let S_i be the set of all mixed strategies of player i.
  - a mixed-strategy profile is an n-tuple $(s_1, \ldots, s_n)$, where $s_i \in S_i$ for each player i.
- Every pure strategy is a mixed strategy.

Expected payoff

Expected payoff

- The goal of each player is to maximize his expected payoff.

Expected payoff

- The goal of each player is to maximize his expected payoff.
- In $G = (P, A, u)$, the expected payoff for player i of the mixed-strategy profile $s = (s_1, \ldots, s_n)$ is

$$u_i(s)$$

Expected payoff

- The goal of each player is to maximize his expected payoff.
- In $G = (P, A, u)$, the expected payoff for player i of the mixed-strategy profile $s = (s_1, \ldots, s_n)$ is

$$u_i(s) = \sum_{a=(a_1,\ldots,a_n)\in A} u_i(a) \cdot \prod_{j=1}^{n} s_j(a_j)$$

Expected payoff

- The goal of each player is to maximize his expected payoff.
- In $G = (P, A, u)$, the expected payoff for player i of the mixed-strategy profile $s = (s_1, \ldots, s_n)$ is

$$u_i(s) = \sum_{a=(a_1,\ldots,a_n)\in A} u_i(a) \cdot \prod_{j=1}^{n} s_j(a_j) = \sum_{\omega\in\Omega} X(\omega) \cdot \Pr[\{\omega\}].$$

Expected payoff

- The goal of each player is to maximize his expected payoff.
- In $G = (P, A, u)$, the expected payoff for player i of the mixed-strategy profile $s = (s_1, \ldots, s_n)$ is

$$u_i(s) = \sum_{a=(a_1,\ldots,a_n)\in A} u_i(a) \cdot \prod_{j=1}^{n} s_j(a_j) = \sum_{\omega\in\Omega} X(\omega) \cdot \Pr[\{\omega\}].$$

  - that is, $u_i(s)$ is the expected value of u_i under the product distribution $\prod_{j=1}^{n} s_j$.

Expected payoff

- The goal of each player is to maximize his expected payoff.
- In $G = (P, A, u)$, the expected payoff for player i of the mixed-strategy profile $s = (s_1, \ldots, s_n)$ is

$$u_i(s) = \sum_{a=(a_1,\ldots,a_n)\in A} u_i(a) \cdot \prod_{j=1}^{n} s_j(a_j) = \sum_{\omega\in\Omega} X(\omega) \cdot \Pr[\{\omega\}].$$

  - that is, $u_i(s)$ is the expected value of u_i under the product distribution $\prod_{j=1}^{n} s_j$.
- It satisfies the linearity of the expected payoff (Exercise):

$$u_i(s) = \sum_{a_i\in A_i} s_i(a_i) \cdot u_i(a_i; s_{-i}),$$

where $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ and
$(s_i'; s_{-i}) = (s_1, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_n)$ for any $s_i' \in S_i$.

Example: expected payoff in Rock-Paper-Scissors

Example: expected payoff in Rock-Paper-Scissors

- Consider the Rock-Paper-Scissors game where each player i uses a strategy s_i that assigns each action the probability $1/3$.

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

Example: expected payoff in Rock-Paper-Scissors

- Consider the Rock-Paper-Scissors game where each player i uses a strategy s_i that assigns each action the probability $1/3$.

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

- By definition, the expected payoff of player 1 on $s = (s_1, s_2)$ is

$$u_1(s) = 1 \cdot \left(\frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3}\right) + (-1) \cdot \left(\frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3}\right)$$
$$+ 0 \cdot \left(\frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3}\right) = 0.$$

Example: expected payoff in Rock-Paper-Scissors

- Consider the Rock-Paper-Scissors game where each player i uses a strategy s_i that assigns each action the probability $1/3$.

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

- By definition, the expected payoff of player 1 on $s = (s_1, s_2)$ is

$$u_1(s) = 1 \cdot \left(\frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3}\right) + (-1) \cdot \left(\frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3}\right)$$
$$+ 0 \cdot \left(\frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3}\right) = 0.$$

- By linearity, $u_1(s) = \frac{1}{3} \cdot 0 + \frac{1}{3} \cdot 0 + \frac{1}{3} \cdot 0 = 0$.

Examples of normal-form games

Examples of normal-form games

- We now give four more examples of normal-form games.

Examples of normal-form games

- We now give four more examples of normal-form games.
- Several of these are used later in the lecture and the tutorials.

Examples of normal-form games

- We now give four more examples of normal-form games.
- Several of these are used later in the lecture and the tutorials.

- We focus here only on two-player games, that is, $P = \{1, 2\}$.

Examples of normal-form games

- We now give four more examples of normal-form games.
- Several of these are used later in the lecture and the tutorials.

- We focus here only on two-player games, that is, $P = \{1, 2\}$.
- These games are called bimatrix games, as they can be represented with two real matrices.

Examples of normal-form games

- We now give four more examples of normal-form games.
- Several of these are used later in the lecture and the tutorials.

- We focus here only on two-player games, that is, $P = \{1, 2\}$.
- These games are called bimatrix games, as they can be represented with two real matrices.
- Player 1 will be the "row player" while player 2 will be the "column player".

Prisoner's dilemma

Prisoner's dilemma

- Two prisoners, are being held in solitary confinement and cannot communicate with the other. Each can either betray the other one by testifying or cooperate with the other one by remaining silent.

Prisoner's dilemma

- Two prisoners, are being held in solitary confinement and cannot communicate with the other. Each can either betray the other one by testifying or cooperate with the other one by remaining silent.

	Testify	Remain silent
Testify	(-2,-2)	(0,-3)
Remain silent	(-3,0)	(-1,-1)

Sources: Serena Maylon (MtG)

Prisoner's dilemma

- Two prisoners, are being held in solitary confinement and cannot communicate with the other. Each can either betray the other one by testifying or cooperate with the other one by remaining silent.

	Testify	Remain silent
Testify	(-2,-2)	(0,-3)
Remain silent	(-3,0)	(-1,-1)

Sources: Serena Maylon (MtG)

- Paradoxically, the only stable solution is when both testify.

Matching pennies

Matching pennies

- We introduce an easier variant of the Rock-Paper-Scissors game.

Matching pennies

- We introduce an easier variant of the Rock-Paper-Scissors game.
- Each of the two players has a penny and chooses either Heads of Tails. If the pennies match, then player 1 wins and keeps both pennies. Otherwise, player 2 keeps both pennies.

Matching pennies

- We introduce an easier variant of the Rock-Paper-Scissors game.
- Each of the two players has a penny and chooses either Heads of Tails. If the pennies match, then player 1 wins and keeps both pennies. Otherwise, player 2 keeps both pennies.

	Heads	Tails
Heads	(1,-1)	(-1,1)
Tails	(-1,1)	(1,-1)

Sources: https://www.fourstateshomepage.com/

Matching pennies

- We introduce an easier variant of the Rock-Paper-Scissors game.
- Each of the two players has a penny and chooses either Heads of Tails. If the pennies match, then player 1 wins and keeps both pennies. Otherwise, player 2 keeps both pennies.

	Heads	Tails
Heads	(1,-1)	(-1,1)
Tails	(-1,1)	(1,-1)

Sources: https://www.fourstateshomepage.com/

- Like Rock-Papers-Scissors, this is a zero-sum game (whatever one player gets, the other one loses).

Matching pennies

- We introduce an easier variant of the Rock-Paper-Scissors game.
- Each of the two players has a penny and chooses either Heads of Tails. If the pennies match, then player 1 wins and keeps both pennies. Otherwise, player 2 keeps both pennies.

	Heads	Tails
Heads	(1,-1)	(-1,1)
Tails	(-1,1)	(1,-1)

Sources: https://www.fourstateshomepage.com/

- Like Rock-Papers-Scissors, this is a zero-sum game (whatever one player gets, the other one loses). Prisoner's dilemma is not.

Battle of sexes

Battle of sexes

- A husband and wife wish to spend an evening together rather than separately, but cannot decide which event to attend. The husband wishes to go to a football match while the wife wants to go to opera.

Battle of sexes

- A husband and wife wish to spend an evening together rather than separately, but cannot decide which event to attend. The husband wishes to go to a football match while the wife wants to go to opera.

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

Sources: https://media.istockphoto.com/

Battle of sexes

- A husband and wife wish to spend an evening together rather than separately, but cannot decide which event to attend. The husband wishes to go to a football match while the wife wants to go to opera.

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

Sources: https://media.istockphoto.com/

- This game displays both cooperation and competition.

Game of chicken

Game of chicken

- Two drivers drive towards each other on a collision course: one must swerve, or both die in the crash. However, if one driver swerves and the other does not, the one who swerved will be called a "chicken"

Game of chicken

- Two drivers drive towards each other on a collision course: one must swerve, or both die in the crash. However, if one driver swerves and the other does not, the one who swerved will be called a "chicken"

	Turn	Go straight
Turn	(0,0)	(-1,1)
Go straight	(1,-1)	(-10,-10)

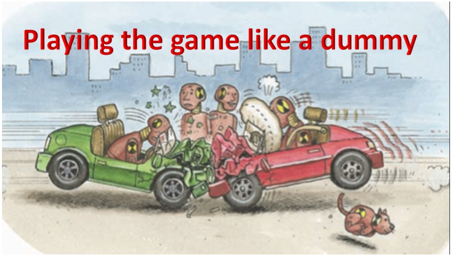

Sources: https://peakd.com/

Game of chicken

- Two drivers drive towards each other on a collision course: one must swerve, or both die in the crash. However, if one driver swerves and the other does not, the one who swerved will be called a "chicken"

	Turn	Go straight
Turn	(0,0)	(-1,1)
Go straight	(1,-1)	(-10,-10)

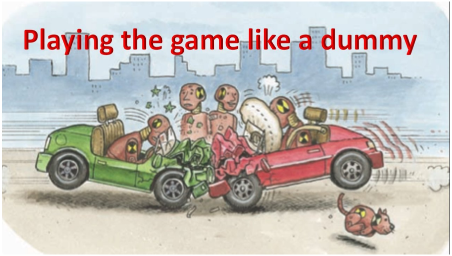

Sources: https://peakd.com/

- What is the best strategy for the players?

Nash equilibrium

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.
- We now introduce perhaps the most influential solution concept, which captures a notion of stability.

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.
- We now introduce perhaps the most influential solution concept, which captures a notion of stability.

- The best response of player i to a strategy profile s_{-i} is a mixed strategy s_i^* such that $u_i(s_i^*; s_{-i}) \geq u_i(s_i'; s_{-i})$ for each $s_i' \in S_i$.

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.
- We now introduce perhaps the most influential solution concept, which captures a notion of stability.

- The best response of player i to a strategy profile s_{-i} is a mixed strategy s_i^* such that $u_i(s_i^*; s_{-i}) \geq u_i(s_i'; s_{-i})$ for each $s_i' \in S_i$.
  - If i knew what strategies the others follow, he would choose this one.

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.
- We now introduce perhaps the most influential solution concept, which captures a notion of stability.

- The best response of player i to a strategy profile s_{-i} is a mixed strategy s_i^* such that $u_i(s_i^*; s_{-i}) \geq u_i(s_i'; s_{-i})$ for each $s_i' \in S_i$.
  - If i knew what strategies the others follow, he would choose this one. It maximizes his expected payoff if others play s_{-i}.

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.
- We now introduce perhaps the most influential solution concept, which captures a notion of stability.

- The best response of player i to a strategy profile s_{-i} is a mixed strategy s_i^* such that $u_i(s_i^*; s_{-i}) \geq u_i(s_i'; s_{-i})$ for each $s_i' \in S_i$.
  - If i knew what strategies the others follow, he would choose this one. It maximizes his expected payoff if others play s_{-i}.
- For a normal-form game $G = (P, A, u)$ of n players, a Nash equilibrium (NE) in G is a strategy profile $(s_1, \ldots, s_n)$ such that s_i is a best response of player i to s_{-i} for every $i \in P$.

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.
- We now introduce perhaps the most influential solution concept, which captures a notion of stability.

- The best response of player i to a strategy profile s_{-i} is a mixed strategy s_i^* such that $u_i(s_i^*; s_{-i}) \geq u_i(s_i'; s_{-i})$ for each $s_i' \in S_i$.
  - If i knew what strategies the others follow, he would choose this one. It maximizes his expected payoff if others play s_{-i}.
- For a normal-form game $G = (P, A, u)$ of n players, a Nash equilibrium (NE) in G is a strategy profile $(s_1, \ldots, s_n)$ such that s_i is a best response of player i to s_{-i} for every $i \in P$.
  - A stable solution concept: no player would like to change his strategy if he knew the strategies of the other players.

Nash equilibrium

- In game theory, we typically study rules for predicting how a game will be played, called solution concepts.
- We now introduce perhaps the most influential solution concept, which captures a notion of stability.

- The best response of player i to a strategy profile s_{-i} is a mixed strategy s_i^* such that $u_i(s_i^*; s_{-i}) \geq u_i(s_i'; s_{-i})$ for each $s_i' \in S_i$.
  - If i knew what strategies the others follow, he would choose this one. It maximizes his expected payoff if others play s_{-i}.
- For a normal-form game $G = (P, A, u)$ of n players, a Nash equilibrium (NE) in G is a strategy profile $(s_1, \ldots, s_n)$ such that s_i is a best response of player i to s_{-i} for every $i \in P$.
  - A stable solution concept: no player would like to change his strategy if he knew the strategies of the other players.
  - Introduced by Nash and by Von Neumann and Morgenstern.

Nash equilibria: remarks

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).
- Example 2: In the Battle of sexes game, there are three Nash equilibria, two pure and one mixed.

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).
- Example 2: In the Battle of sexes game, there are three Nash equilibria, two pure and one mixed.

- Do Nash equilibria always exist in every game?

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).
- Example 2: In the Battle of sexes game, there are three Nash equilibria, two pure and one mixed.

- Do Nash equilibria always exist in every game? Is there always a stable solution concept?

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).
- Example 2: In the Battle of sexes game, there are three Nash equilibria, two pure and one mixed.

- Do Nash equilibria always exist in every game? Is there always a stable solution concept?
- Yes, they do!

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).
- Example 2: In the Battle of sexes game, there are three Nash equilibria, two pure and one mixed.

- Do Nash equilibria always exist in every game? Is there always a stable solution concept?
- Yes, they do! Shown by Nash in 1950.

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).
- Example 2: In the Battle of sexes game, there are three Nash equilibria, two pure and one mixed.

- Do Nash equilibria always exist in every game? Is there always a stable solution concept?
- Yes, they do! Shown by Nash in 1950.
- Maybe the most influential result in game theory.

Nash equilibria: remarks

- Neither best responses nor Nash equilibria are determined uniquely.
- Example 1: In the Rock-Paper-Scissors game, there is a unique mixed Nash equilibrium (both players play everything with probability $1/3$).
- Example 2: In the Battle of sexes game, there are three Nash equilibria, two pure and one mixed.

- Do Nash equilibria always exist in every game? Is there always a stable solution concept?
- Yes, they do! Shown by Nash in 1950.
- Maybe the most influential result in game theory. Later, Nash received a Nobel prize for economics.

Nash's Theorem

Nash's Theorem

Nash's Theorem (Theorem 2.16)

Every normal-form game has a Nash equilibrium.

Nash's Theorem

Nash's Theorem (Theorem 2.16)

Every normal-form game has a Nash equilibrium.

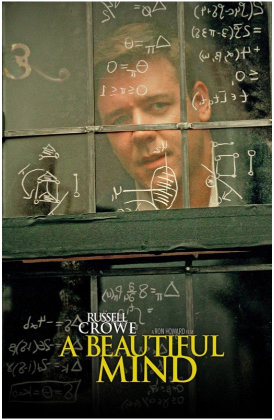

Figure: John Forbes Nash Jr. (1928–2015) and his depiction in the movie A Beautiful mind.

Sources: https://britannica.com and https://medium.com

Preparations for the proof of Nash's theorem

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem.

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

- For $d \in \mathbb{N}$, a subset X of $\mathbb{R}^d$ is compact if X is closed and bounded.

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

- For $d \in \mathbb{N}$, a subset X of $\mathbb{R}^d$ is compact if X is closed and bounded.
- We say that a subset Y of $\mathbb{R}^d$ is convex if every line segment containing two points from Y is fully contained in Y.

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

- For $d \in \mathbb{N}$, a subset X of $\mathbb{R}^d$ is compact if X is closed and bounded.
- We say that a subset Y of $\mathbb{R}^d$ is convex if every line segment containing two points from Y is fully contained in Y. Formally: for all x, y from Y, $tx + (1-t)y \in Y$ for every $t \in [0,1]$.

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

- For $d \in \mathbb{N}$, a subset X of $\mathbb{R}^d$ is compact if X is closed and bounded.
- We say that a subset Y of $\mathbb{R}^d$ is convex if every line segment containing two points from Y is fully contained in Y. Formally: for all x, y from Y, $tx + (1-t)y \in Y$ for every $t \in [0,1]$.
- For n affinely independent points $x_1, \dots, x_n \in \mathbb{R}^d$, an $(n-1)$-simplex Δ_n on $x_1, \dots, x_n$ is the set of convex combinations of the points $x_1, \dots, x_n$.

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

- For $d \in \mathbb{N}$, a subset X of $\mathbb{R}^d$ is compact if X is closed and bounded.
- We say that a subset Y of $\mathbb{R}^d$ is convex if every line segment containing two points from Y is fully contained in Y. Formally: for all x, y from Y, $tx + (1-t)y \in Y$ for every $t \in [0,1]$.
- For n affinely independent points $x_1, \dots, x_n \in \mathbb{R}^d$, an $(n-1)$-simplex Δ_n on $x_1, \dots, x_n$ is the set of convex combinations of the points $x_1, \dots, x_n$. Each simplex is a compact convex set in $\mathbb{R}^d$.

Preparations for the proof of Nash's theorem

- The proof is essentially topological, as its main ingredient is a fixed-point theorem. We use a theorem due to Brouwer.

- For $d \in \mathbb{N}$, a subset X of $\mathbb{R}^d$ is compact if X is closed and bounded.
- We say that a subset Y of $\mathbb{R}^d$ is convex if every line segment containing two points from Y is fully contained in Y. Formally: for all x, y from Y, $tx + (1-t)y \in Y$ for every $t \in [0,1]$.
- For n affinely independent points $x_1, \ldots, x_n \in \mathbb{R}^d$, an $(n-1)$-simplex Δ_n on $x_1, \ldots, x_n$ is the set of convex combinations of the points $x_1, \ldots, x_n$. Each simplex is a compact convex set in $\mathbb{R}^d$.

Lemma (Lemma 2.18)

For $n, d_1, \ldots, d_n \in \mathbb{N}$, let $K_1, \ldots, K_n$ be compact sets, each K_i lying in $\mathbb{R}^{d_i}$. Then, $K_1 \times \cdots \times K_n$ is a compact set in $\mathbb{R}^{d_1+\cdots+d_n}$.

Brouwer's Fixed Point Theorem

Brouwer's Fixed Point Theorem

Brouwer's Fixed Point Theorem (Thoerem 2.17)

For each $d \in \mathbb{N}$, let K be a non-empty compact convex set in $\mathbb{R}^d$ and $f : K \to K$ be a continuous mapping. Then, there exists a fixed point $x_0 \in K$ for f, that is, $f(x_0) = x_0$.

Brouwer's Fixed Point Theorem

Brouwer's Fixed Point Theorem (Thoerem 2.17)

For each $d \in \mathbb{N}$, let K be a non-empty compact convex set in $\mathbb{R}^d$ and $f : K \to K$ be a continuous mapping. Then, there exists a fixed point $x_0 \in K$ for f, that is, $f(x_0) = x_0$.

Figure: L. E. J. Brouwer (1881–1966).

Source: https://arxiv.org/pdf/1612.06820.pdf

Figure: John Forbes Nash Jr. receiving a Nobel prize for economics.

Source: https://pbs.org

Figure: John Forbes Nash Jr. receiving a Nobel prize for economics.

Source: https://pbs.org

Thank you for your attention.