Algorithmic game theory

Martin Balko

3rd lecture

October 19th 2023

Proof of the Minimax Theorem

The Minimax Theorem

The Minimax Theorem

- For every zero-sum game, worst-case optimal strategies for both players exist and can be efficiently computed. There is a number v such that, for any worst-case optimal strategies x^* and y^*, the strategy profile (x^*, y^*) is a Nash equilibrium and $\beta(x^*) = (x^*)^\top My^* = \alpha(y^*) = v$.

The Minimax Theorem

- For every zero-sum game, worst-case optimal strategies for both players exist and can be efficiently computed. There is a number v such that, for any worst-case optimal strategies x^* and y^*, the strategy profile (x^*, y^*) is a Nash equilibrium and $\beta(x^*) = (x^*)^\top M y^* = \alpha(y^*) = v$.

Figure: John von Neumann (1903–1957) and Oskar Morgenstern (1902–1977).

Sources: https://en.wikiquote.org and https://austriainusa.org

The Minimax Theorem

- For every zero-sum game, worst-case optimal strategies for both players exist and can be efficiently computed. There is a number v such that, for any worst-case optimal strategies x^* and y^*, the strategy profile (x^*, y^*) is a Nash equilibrium and $\beta(x^*) = (x^*)^\top M y^* = \alpha(y^*) = v$.

The Minimax Theorem

- For every zero-sum game, worst-case optimal strategies for both players exist and can be efficiently computed. There is a number v such that, for any worst-case optimal strategies x^* and y^*, the strategy profile (x^*, y^*) is a Nash equilibrium and $\beta(x^*) = (x^*)^\top M y^* = \alpha(y^*) = v$.

- Recall that $\beta(x) = \min_{y \in S_2} x^\top My$ and $\alpha(y) = \max_{x \in S_1} x^\top My$ are the best possible payoffs of player 2 to x and of player 1 to y, respectively.

The Minimax Theorem

- For every zero-sum game, worst-case optimal strategies for both players exist and can be efficiently computed. There is a number v such that, for any worst-case optimal strategies x^* and y^*, the strategy profile (x^*, y^*) is a Nash equilibrium and $\beta(x^*) = (x^*)^\top M y^* = \alpha(y^*) = v$.

- Recall that $\beta(x) = \min_{y \in S_2} x^\top M y$ and $\alpha(y) = \max_{x \in S_1} x^\top M y$ are the best possible payoffs of player 2 to x and of player 1 to y, respectively.
- Also, the worst-case optimal strategy $\overline{x}$ for player 1, satisfies

$$\beta(\overline{x}) = \max_{x \in S_1} \beta(x).$$

The Minimax Theorem

- For every zero-sum game, worst-case optimal strategies for both players exist and can be efficiently computed. There is a number v such that, for any worst-case optimal strategies x^* and y^*, the strategy profile (x^*, y^*) is a Nash equilibrium and $\beta(x^*) = (x^*)^\top M y^* = \alpha(y^*) = v$.

- Recall that $\beta(x) = \min_{y \in S_2} x^\top M y$ and $\alpha(y) = \max_{x \in S_1} x^\top M y$ are the best possible payoffs of player 2 to x and of player 1 to y, respectively.
- Also, the worst-case optimal strategy $\bar{x}$ for player 1, satisfies

$$\beta(\bar{x}) = \max_{x \in S_1} \beta(x).$$

- The worst-case optimal strategy $\bar{y}$ for player 2, satisfies

$$\alpha(\bar{y}) = \min_{y \in S_2} \alpha(y).$$

Duality of linear programming

Duality of linear programming

	Primal linear program	Dual linear program
Variables	$x_1, \ldots, x_m$	$y_1, \ldots, y_n$
Matrix	$A \in \mathbb{R}^{n \times m}$	$A^\top \in \mathbb{R}^{m \times n}$
Right-hand side	$b \in \mathbb{R}^n$	$c \in \mathbb{R}^m$
Objective function	$\max c^\top x$	$\min b^\top y$
Constraints	ith constraint has $\leq$	$y_i \geq 0$
	$\geq$	$y_i \leq 0$
	$=$	$y_i \in \mathbb{R}$
	$x_j \geq 0$	jth constraint has $\geq$
	$x_j \leq 0$	$\leq$
	$x_j \in \mathbb{R}$	$=$

Table: A recipe for making dual programs.

Nash equilibria in bimatrix games

Bimatrix games

Bimatrix games

- We try to efficiently find Nash equilibria in bimatrix games, that is, games of 2-players (not necessarily zero-sum).

Bimatrix games

- We try to efficiently find Nash equilibria in bimatrix games, that is, games of 2-players (not necessarily zero-sum).
- Example: Prisoner's dilemma

Bimatrix games

- We try to efficiently find Nash equilibria in bimatrix games, that is, games of 2-players (not necessarily zero-sum).
- Example: Prisoner's dilemma

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

Sources: https://sciworthy.com/

Bimatrix games examples: collaborative projects

Source: https://filestage.io/

Bimatrix games examples: education, knowledge sharing

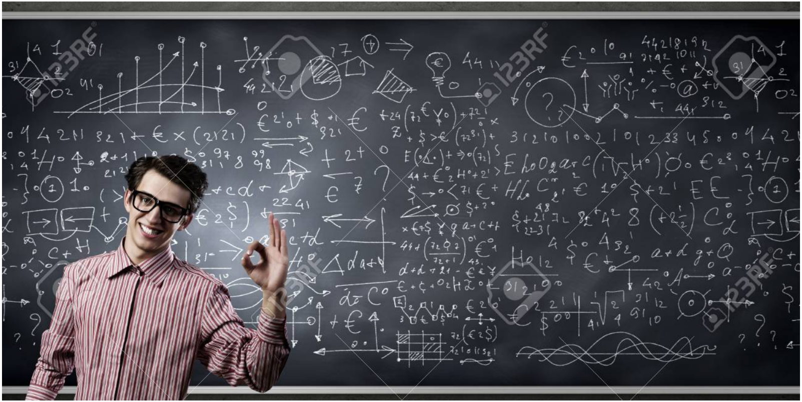

Source: https://www.123rf.com/

Bimatrix games examples: the battle for Gotham's soul

	Cooperate	Detonate
Cooperate	(0,0)	(0,1)
Detonate	(1,0)	(0,0)

Sources: https://www.cbr.com/

Nash equilibria in bimatrix games by brute force

Nash equilibria in bimatrix games by brute force

- We try to design an algorithm for finding Nash equilibria in games of two players (bimatrix games).

Nash equilibria in bimatrix games by brute force

- We try to design an algorithm for finding Nash equilibria in games of two players (bimatrix games).
- We state some observations that yield a brute-force algorithm.

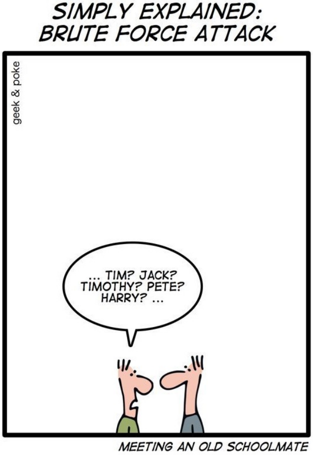

Source: https://pinterest.com

Nash equilibria in bimatrix games by brute force

- We try to design an algorithm for finding Nash equilibria in games of two players (bimatrix games).
- We state some observations that yield a brute-force algorithm.

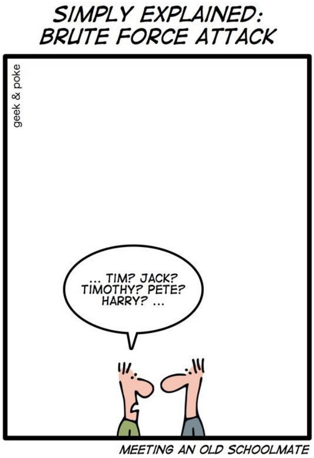

Source: https://pinterest.com

- Later, we show the currently best known algorithm for this problem.

Example: Battle of sexes

Example: Battle of sexes

- We show the brute-force algorithm on the Battle of sexes game.

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

Example: Battle of sexes

- We show the brute-force algorithm on the Battle of sexes game.

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

- That is, we have $M = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$ and $N = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} = N^\top$.

Example: Battle of sexes

- We show the brute-force algorithm on the Battle of sexes game.

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

- That is, we have $M = \left(\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ and $N = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right) = N^\top$.
- If $I = \{1, 2\}$ and $J = \{1, 2\}$, then we want to solve the following system of 6 equations with 6 variables x_1, x_2, y_1, y_2, u, v:

$$x_1 = v, \quad 2x_2 = v, \quad x_1 + x_2 = 1$$
$$2y_1 = u, \quad y_2 = u, \; ; \; y_1 + y_2 = 1$$

Example: Battle of sexes

- We show the brute-force algorithm on the Battle of sexes game.

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

- That is, we have $M = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$ and $N = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} = N^\top$.
- If $I = \{1, 2\}$ and $J = \{1, 2\}$, then we want to solve the following system of 6 equations with 6 variables x_1, x_2, y_1, y_2, u, v:

$$x_1 = v, \quad 2x_2 = v, \quad x_1 + x_2 = 1$$
$$2y_1 = u, \quad y_2 = u, \; ; \; y_1 + y_2 = 1$$

- This yields a unique solution $(x_1, x_2) = (\frac{2}{3}, \frac{1}{3})$ and $(y_1, y_2) = (\frac{1}{3}, \frac{2}{3})$.

Example: Battle of sexes

- We show the brute-force algorithm on the Battle of sexes game.

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

- That is, we have $M = \left(\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ and $N = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right) = N^\top$.
- If $I = \{1, 2\}$ and $J = \{1, 2\}$, then we want to solve the following system of 6 equations with 6 variables x_1, x_2, y_1, y_2, u, v:

$$x_1 = v, \quad 2x_2 = v, \quad x_1 + x_2 = 1$$
$$2y_1 = u, \quad y_2 = u, \; ; \; y_1 + y_2 = 1$$

- This yields a unique solution $(x_1, x_2) = (\frac{2}{3}, \frac{1}{3})$ and $(y_1, y_2) = (\frac{1}{3}, \frac{2}{3})$. Since $x, y \geq \mathbf{0}$ and there is no better pure strategy, we have NE.

Preliminaries from geometry

Preliminaries from geometry

- A hyperplane in $\mathbb{R}^d$ is a set $\{x \in \mathbb{R}^d \colon v^\top x = w\}$ for some $v \in \mathbb{R}^d$ and $w \in \mathbb{R}$.

Preliminaries from geometry

- A hyperplane in $\mathbb{R}^d$ is a set $\{x \in \mathbb{R}^d \colon v^\top x = w\}$ for some $v \in \mathbb{R}^d$ and $w \in \mathbb{R}$.
- A halfspace in $\mathbb{R}^d$ is a set $\{x \in \mathbb{R}^d \colon v^\top x \leq w\}$.

Preliminaries from geometry

- A hyperplane in $\mathbb{R}^d$ is a set $\{x \in \mathbb{R}^d \colon v^\top x = w\}$ for some $v \in \mathbb{R}^d$ and $w \in \mathbb{R}$.
- A halfspace in $\mathbb{R}^d$ is a set $\{x \in \mathbb{R}^d \colon v^\top x \leq w\}$.

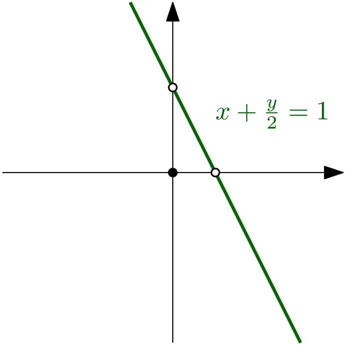

Preliminaries from geometry

- A hyperplane in $\mathbb{R}^d$ is a set $\{x \in \mathbb{R}^d \colon v^\top x = w\}$ for some $v \in \mathbb{R}^d$ and $w \in \mathbb{R}$.
- A halfspace in $\mathbb{R}^d$ is a set $\{x \in \mathbb{R}^d \colon v^\top x \leq w\}$.

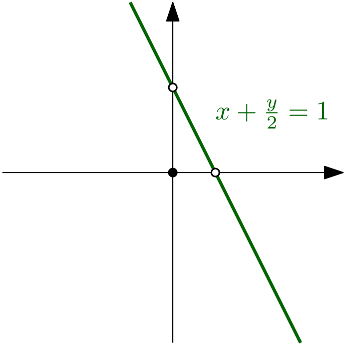

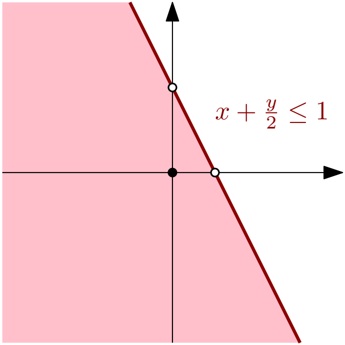

Preliminaries from geometry

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$.

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$. That is, $P = \{x \in \mathbb{R}^d \colon Vx \leq u\}$ for some $V \in \mathbb{R}^{n \times d}$ and $u \in \mathbb{R}^n$, where n is the number of halfspaces determining P.

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$. That is, $P = \{x \in \mathbb{R}^d \colon Vx \leq u\}$ for some $V \in \mathbb{R}^{n \times d}$ and $u \in \mathbb{R}^n$, where n is the number of halfspaces determining P.
- A bounded polyhedron is called polytope.

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$. That is, $P = \{x \in \mathbb{R}^d : Vx \leq u\}$ for some $V \in \mathbb{R}^{n \times d}$ and $u \in \mathbb{R}^n$, where n is the number of halfspaces determining P.
- A bounded polyhedron is called polytope.

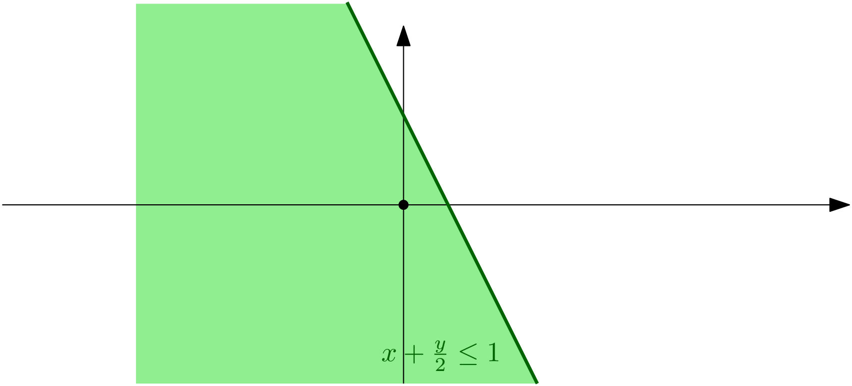

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$. That is, $P = \{x \in \mathbb{R}^d \colon Vx \leq u\}$ for some $V \in \mathbb{R}^{n \times d}$ and $u \in \mathbb{R}^n$, where n is the number of halfspaces determining P.
- A bounded polyhedron is called polytope.

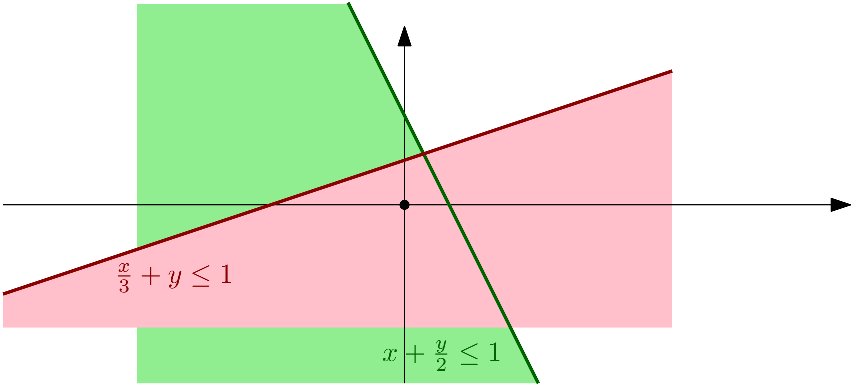

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$. That is, $P = \{x \in \mathbb{R}^d : Vx \leq u\}$ for some $V \in \mathbb{R}^{n \times d}$ and $u \in \mathbb{R}^n$, where n is the number of halfspaces determining P.
- A bounded polyhedron is called polytope.

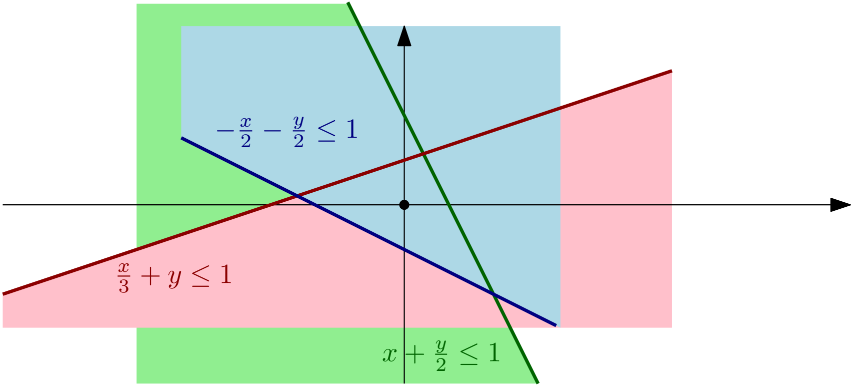

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$. That is, $P = \{x \in \mathbb{R}^d : Vx \leq u\}$ for some $V \in \mathbb{R}^{n \times d}$ and $u \in \mathbb{R}^n$, where n is the number of halfspaces determining P.
- A bounded polyhedron is called polytope.

$$V = \begin{pmatrix} 1 & \frac{1}{2} \\ \frac{1}{3} & 1 \\ -\frac{1}{2} & -\frac{1}{2} \end{pmatrix}$$

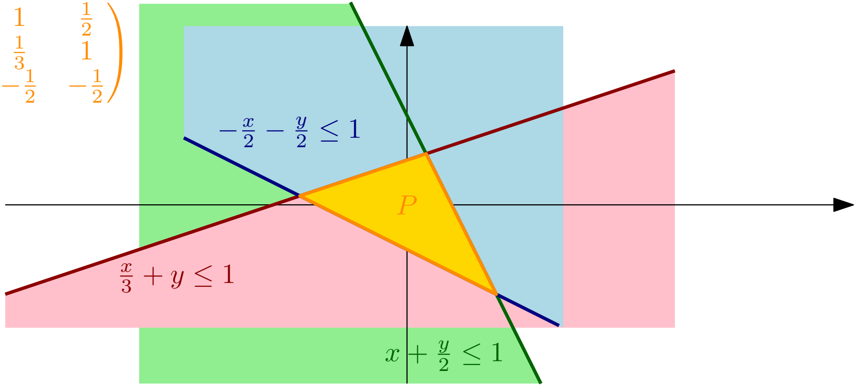

Preliminaries from geometry

- A (convex) polyhedron P in $\mathbb{R}^d$ is an intersection of finitely many halfspaces in $\mathbb{R}^d$. That is, $P = \{x \in \mathbb{R}^d : Vx \leq u\}$ for some $V \in \mathbb{R}^{n \times d}$ and $u \in \mathbb{R}^n$, where n is the number of halfspaces determining P.
- A bounded polyhedron is called polytope.

$$V = \begin{pmatrix} 1 & \frac{1}{2} \\ \frac{1}{3} & 1 \\ -\frac{1}{2} & -\frac{1}{2} \end{pmatrix}$$

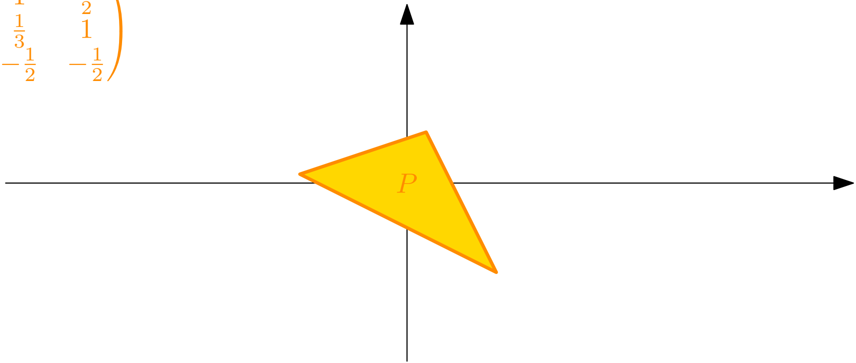

Examples of polytopes in $\mathbb{R}^3$

Examples of polytopes in $\mathbb{R}^3$

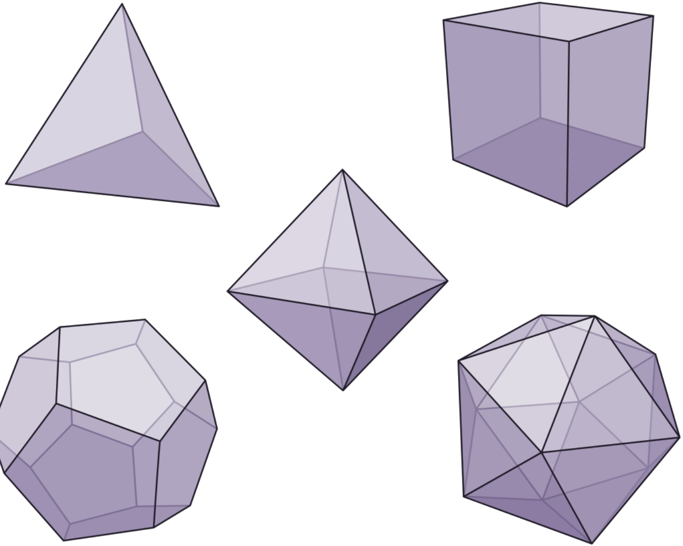

Best response polyhedra $\overline{P}$ and $\overline{Q}$ for the Battle of sexes

Best response polyhedra $\overline{P}$ and $\overline{Q}$ for the Battle of sexes

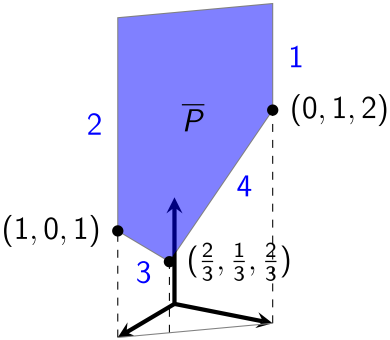

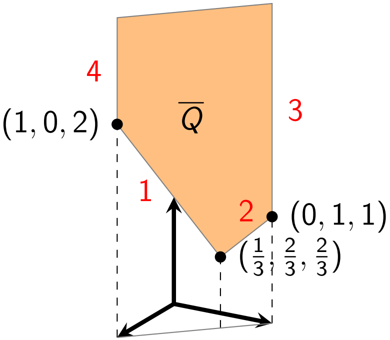

$$\overline{P} = \{(x_1, x_2, v) \in \mathbb{R}^2 \times \mathbb{R} \colon x_1, x_2 \geq \mathbf{0}, x_1 + x_2 = 1, x_1 \leq v, 2x_2 \leq v\}$$

$$\overline{Q} = \{(y_3, y_4, u) \in \mathbb{R}^2 \times \mathbb{R} \colon y_3, y_4 \geq \mathbf{0}, y_3 + y_4 = 1, 2y_3 \leq u, y_4 \leq u\}.$$

- Next lecture we learn the Lemke–Howson algorithm, the best known algorithm to find Nash equilibria in bimatrix games.

- Next lecture we learn the Lemke–Howson algorithm, the best known algorithm to find Nash equilibria in bimatrix games.

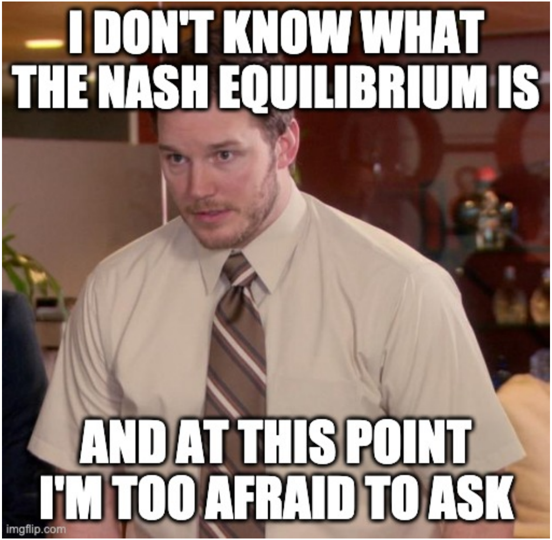

- Next lecture we learn the Lemke–Howson algorithm, the best known algorithm to find Nash equilibria in bimatrix games.

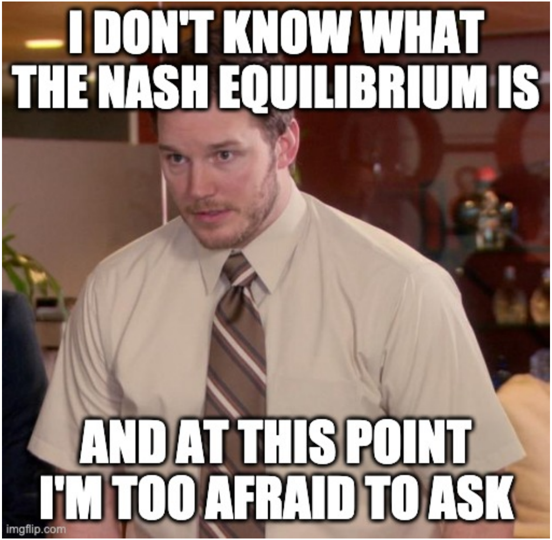

Thank you for your attention.