An Introduction to Non-Cooperative Game Theory

Siddharth Barman

Indian Institute of Science

Game Theory: Study of how self-interested agents interact.

Game Theory: Study of how self-interested agents interact.

This field uses mathematical models to understand and predict outcomes of strategic interactions.

Game Theory: Study of how self-interested agents interact.

This field uses mathematical models to understand and predict outcomes of strategic interactions.

Game Components: Players, Actions, Payoffs

Game Theory: Study of how self-interested agents interact.

This field uses mathematical models to understand and predict outcomes of strategic interactions.

Game Components: Players, Actions, Payoffs

Likely/Stable Outcome: *Equilibria* wherein no player has an incentive to unilaterally deviate

Game Theory: Study of how self-interested agents interact.

Illustrative Applications:

Game Theory: Study of how self-interested agents interact.

Illustrative Applications:

- Traffic Networks: Self-interested users strategically choosing routes in a network to minimize the delay they face.

Insight: Formal explanation of Braess' paradox

Game Theory: Study of how self-interested agents interact.

Illustrative Applications:

- Traffic Networks
- Auctions: Strategic vendor auctioning goods to self-interested bidders

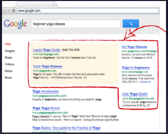

Insight: A simple auction with one extra bidder earns more revenue than the optimal auction with the original bidders (Bulow and Klemperer 1996).

Game Theory: Study of how self-interested agents interact.

Illustrative Applications:

- Traffic Networks
- Auctions
- Stable Matchings: Determine a stable assignment for self-interested entities that have rankings for each other

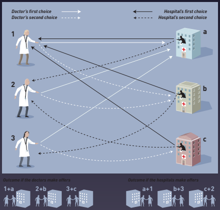

Insight: The stark effect of competition (Ashlagi et al. 2015).

Game Theory: Study of how self-interested agents interact.

Game Components: Players, Actions, Payoffs

Likely/Stable Outcome: *Equilibria* wherein no player has an incentive to unilaterally deviate

Game Theory: Study of how self-interested agents interact.

Components of a Game: Players, Actions, Payoffs

Likely/Stable Outcome: *Equilibria* wherein no player has an incentive to unilaterally deviate

Representation of a Game:

- Normal Form
- ~~Extensive Form~~

Game Theory: Study of how self-interested agents interact.

Components of a Game: Players, Actions, Payoffs

Likely/Stable Outcome: *Equilibria* wherein no player has an incentive to unilaterally deviate

Representation of a Game:

- Normal Form includes all action profiles and their corresponding payoffs, for each player

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

(me circa 1990)

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

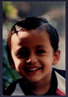

(me circa 1990)

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

At (E, A) no player has an incentive to unilaterally deviate

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

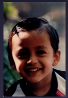

(me circa 1990)

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

At (E, A) and at (NE, NA) no player has an incentive to unilaterally deviate

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Presentation Game[1]

	Pay attention (A)	Do not pay attention (NA)
Put effort into presentation (E)	2, 2	-1, 0
Do not put effort into presentation (NE)	-7, -8	0, 0

(me circa 1990)

(E, A) and (NE, NA) are Pure Nash Equilibria of the game

[1] Credit: Vincent Conitzer

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Rock-Paper-Scissors

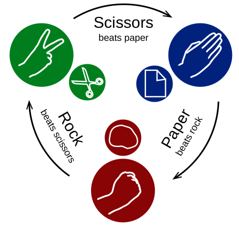

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Notation:

$$u_1(R, P) = -1$$

$$u_2(R, P) = 1$$

....

Two-Player Games model settings in which two self-interested entities *simultaneously* select actions to maximize their own payoffs.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Amongst rational players, deterministic strategies are <u>not</u> stable.

Therefore, we must consider strategies in which players randomize between actions.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Notation:

$$u_1(R, P) = -1$$

$$u_2(R, P) = 1$$

....

This is a *zero-sum* game

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\text{R, P, S}\}$.
- Expected utility of first player
  $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\mathrm{R, P, S}\}$.
- Expected utility of first player
  $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$. Hence, $u_1(\sigma, \sigma) = 0$.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\mathrm{R, P, S}\}$.
- Expected utility of first player $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$. Hence, $u_1(\sigma, \sigma) = 0$.
- Also, $u_2(\sigma, R) = u_2(\sigma, P) = u_2(\sigma, S) = u_2(\sigma, \sigma) = 0$.

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\text{R, P, S}\}$.
- Expected utility of first player $u_1(R, \sigma) = u_1(P, \sigma) = u_1(S, \sigma) = 0$. Hence, $u_1(\sigma, \sigma) = 0$.
- Also, $u_2(\sigma, R) = u_2(\sigma, P) = u_2(\sigma, S) = u_2(\sigma, \sigma) = 0$.

At (σ, σ) players have no incentive to unilaterally deviate

Example: Rock-Paper-Scissors

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

- $\sigma :=$ uniform distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over $\{\text{R, P, S}\}$.
- Expected utility of first player $u_1(R,\sigma) = u_1(P,\sigma) = u_1(S,\sigma) = 0$. Hence, $u_1(\sigma,\sigma) = 0$.
- Also, $u_2(\sigma,R) = u_2(\sigma,P) = u_2(\sigma,S) = u_2(\sigma,\sigma) = 0$.

(σ,σ) is a Nash equilibrium of the game

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Players: $\{1, 2, \ldots, n\}$
- Player p's action set: A_p

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Players: $\{1, 2, \ldots, n\}$
- Player p's action set: A_p
- Action profiles: $\prod_p A_p$
- Player p's utility: $u_p : \prod_p A_p \to \mathbb{R}$

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Players: $\{1, 2, \ldots, n\}$
- Player p's action set: A_p
- Action profiles: $\prod_p A_p$
- Player p's utility: $u_p : \prod_p A_p \to \mathbb{R}$

In Rock-Paper-Scissors, $n = 2$ and $A_1 = A_2 = \{R, P, S\}$
$u_1(R, P) = -1$, $u_2(R, P) = 1$,...

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Notation:

- Players: $\{1, 2, \ldots, n\}$
- Player p's action set: A_p
- Action profiles: $\prod_p A_p$
- Player p's utility: $u_p : \prod_p A_p \to \mathbb{R}$

Probability distributions $(\sigma_1, \sigma_2, \ldots, \sigma_n)$ denote a Nash equilibrium iff for each player p we have

$$u_p(a_p, \sigma_{-p}) \leq u_p(\sigma_p, \sigma_{-p}) \qquad \forall a_p \in A_p.$$

Here, $\sigma_{-p} := (\sigma_1, \sigma_2, \ldots, \sigma_{p-1}, \sigma_{p+1}, \ldots, \sigma_n)$.

Fundamental Results

Guaranteed Existence of Nash Equilibria

- In two-player zero-sum games [von Neumann 1928]
- In finite games [Nash 1950]

John von Neumann

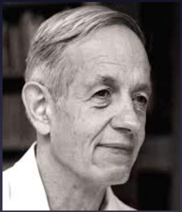

John Nash

Two-Player Zero-Sum Games

Recall Rock-Paper-Scissors:

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Two-Player Zero-Sum Games

In general, for each action $a_1 \in A_1$ and $a_2 \in A_2$

$$u_1(a_1, a_2) + u_2(a_1, a_2) = 0$$

Two-Player Zero-Sum Games

In general, for any action $a_1 \in A_1$ and $a_2 \in A_2$

$$u_2(a_1, a_2) = -u_1(a_1, a_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\max_{\sigma_1 \in \Delta(A_1)} \min_{\sigma_2 \in \Delta(A_2)} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value $=$ largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value $=$ largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value $=$ smallest utility that 2 can force on 1

$$\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

Maximin value = Minimax value

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\max_{\sigma_1} \min_{\sigma_2} \; u_1(\sigma_1, \sigma_2) = \min_{\sigma_1} \max_{\sigma_1} \; u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Minimax Theorem $\Rightarrow$ Existence of Nash Eq. in zero-sum games

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

$$u_1(\sigma_1^*, \sigma_2^*) \geq \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2)$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min\nolimits_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max\nolimits_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

$$\begin{aligned} u_1(\sigma_1^*, \sigma_2^*) &\geq \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \\ &= \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \end{aligned}$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

$$\begin{aligned} u_1(\sigma_1^*, \sigma_2^*) &\geq \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) \\ &= \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*) \\ &\geq u_1(a_1, \sigma_2^*) \qquad \forall a_1 \in A_1 \end{aligned}$$

Two-Player Zero-Sum Games

- Maximin value = largest utility that player 1 can guarantee

$$\sigma_1^* \in \arg\max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

- Minimax value = smallest utility that 2 can force on 1

$$\sigma_2^* \in \arg\min_{\sigma_2} \max_{\sigma_1} u_1(\sigma_1, \sigma_2)$$

Minimax Theorem (von Neumann 1928)

$$\min_{\sigma_2} u_1(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Minimax Theorem ⇒ Existence of Nash Eq. in zero-sum games □

Fundamental Results

Guaranteed Existence of Nash Equilibria

✓ In two-player zero-sum games [von Neumann 1928]

- In finite games [Nash 1950]


John von Neumann

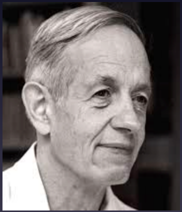
John Nash

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

- Players $\{1, 2, \ldots, n\}$ and action sets $A_1, A_2, \ldots, A_n$
- Player p's utility: $u_p : \prod_p A_p \to \mathbb{R}$

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

- Players $\{1, 2, \ldots, n\}$ and action sets $A_1, A_2, \ldots, A_n$
- Player p's utility: $u_p : \prod_p A_p \to \mathbb{R}$

Prob. dist. $(\sigma_1, \sigma_2, \ldots, \sigma_n)$ denote a Nash equilibrium iff

$$u_p(a_p, \sigma_{-p}) \leq u_p(\sigma_p, \sigma_{-p}) \qquad \forall p,\ \forall a_p \in A_p.$$

Here, $\sigma_{-p} := (\sigma_1, \sigma_2, \ldots, \sigma_{p-1}, \sigma_{p+1}, \ldots, \sigma_n)$.

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Nash's Existence Theorem (1950)

Every finite game has at least one Nash equilibrium.

Nash equilibria denote distributions over players' action profiles at which no player can benefit by unilateral deviation.

Nash's Existence Theorem (1950)

Every finite game has at least one Nash equilibrium.

Proof via Brouwer's fixed point theorem.

Fundamental Results

Guaranteed Existence of Nash Equilibria

- ✓ In two-player zero-sum games [von Neumann 1928]
- ✓ In finite games [Nash 1950]

John von Neumann

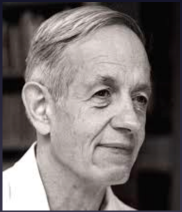

John Nash

Algorithmic Game Theory

Algorithmic Game Theory

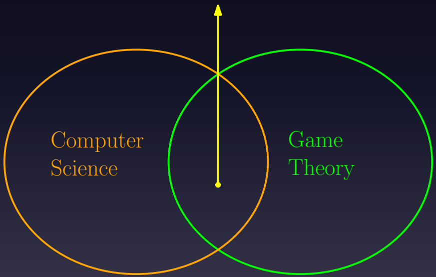

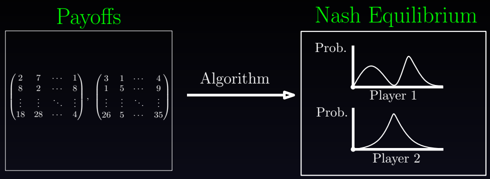

Payoffs
$\begin{pmatrix} 2 & 7 & \cdots & 1 \\ 8 & 2 & \cdots & 8 \\ \vdots & \vdots & \ddots & \vdots \\ 18 & 28 & \cdots & 4 \end{pmatrix}, \begin{pmatrix} 3 & 1 & \cdots & 4 \\ 1 & 5 & \cdots & 9 \\ \vdots & \vdots & \ddots & \vdots \\ 26 & 5 & \cdots & 35 \end{pmatrix}$
Algorithm
Nash Equilibrium
Prob.
Player 1
Prob.
Player 2

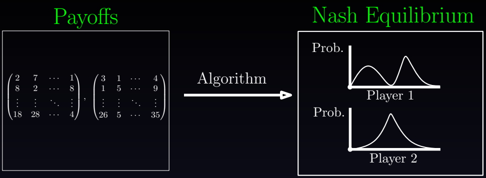

> Nash equilibria of zero-sum games can be computed in polynomial time.

Minimax strategies via linear programming [Dantzig 1951].

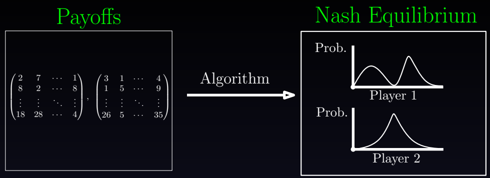

Complexity of Equilibria

✓ Zero-Sum Games

- General Two-Player Games?
- Multi-Player Games?

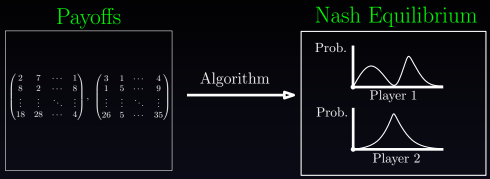

Every instance of NASH admits a solution

NP-hardness cannot be applied to such problems

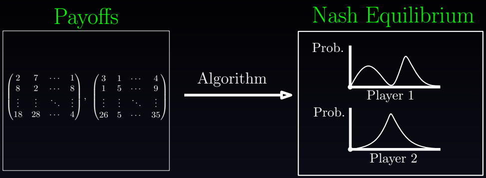

$$\textsc{Nash} \in \text{PPAD}$$

PPAD (Polynomial Parity Arguments on Directed graphs) := Probs. that can be solved via directed path-following algorithms.

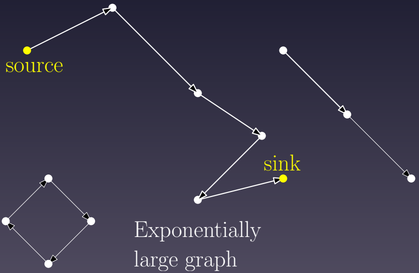

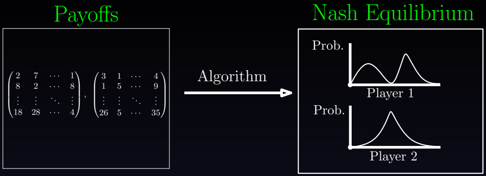

$$\textsc{Nash} \in \text{PPAD}$$

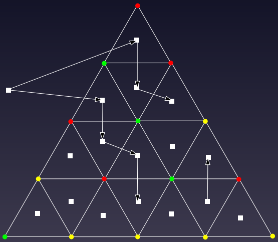

Sperner's Lemma

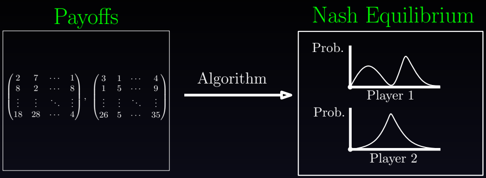

NASH is PPAD-complete

Even for two player games [DGP06, CDT09]

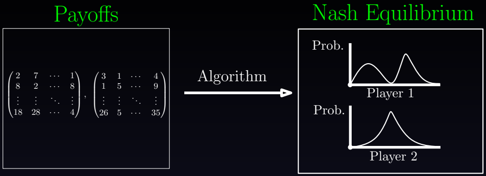

NASH is PPAD-hard

Even for two player games [DGP06, CDT09]

Central Open Question: A polynomial-time algorithm for *approximate* Nash?

Additional Topics

- Extensive-form games
- Equilibrium refinements
- Games with imperfect information
- No-regret dynamics
- Other solution concepts, e.g. correlated eq.
- ...

Selected References:

Game Theory: Analysis of Conflict. R.B. Myerson

A Course in Game Theory. M.J. Osborne and A. Rubinstein

Game Theory and Mechanism Design. Y. Narahari

Algorithmic Game Theory. Nisan et al.

Additional Topics

- Extensive-form games
- Equilibrium refinements
- Games with imperfect information
- No-regret dynamics
- Other solution concepts, e.g. correlated eq.
- ...

Selected References:

Game Theory: Analysis of Conflict. R.B. Myerson

A Course in Game Theory. M.J. Osborne and A. Rubinstein

Game Theory and Mechanism Design. Y. Narahari

Algorithmic Game Theory. Nisan et al.

Thank You!