

Algorithmic Game Theory

CS 58100 · Game Theory I

Alex Psomas

Fall 2026

HOW DO PEOPLE MAKE DECISIONS?

- At the core of the course is the concept of a rational agent.
- In this lecture, we will discuss a core theorem in economics, that will justify expected utility maximization, the main tool/assumption we'll use throughout the course

OUTCOMES AND LOTTERIES

- Let $O = \{o_1, \dots, o_n\}$ be a set of outcomes
  - For example, $O =$ {"Everyone gets an A+ in CS581", "The Odyssey wins every single Oscar" , "NVIDIA stock is worth 1000$ by end-of-year"}
- A lottery is a distribution over outcomes
  - For example, "get 1000$ w.p. ½, get a hat w.p. 1/3, get nothing w.p. 1/6"
- We write $[p_1 : o_1, p_2 : o_2, \dots, p_n : o_n]$ for the lottery that assigns probability p_1 to outcome o_1, and so on.
- We write $\Delta(O)$ for the set of lotteries on O

ORDINAL PREFERENCES

- A preference relation $\succcurlyeq$ indicates the relative preference of an agent for outcomes and lotteries
- We write:
  - $o_1 \succcurlyeq o_2$ if the agent weakly prefers outcome o_1 to outcome o_2
  - $o_1 \succ o_2$ if the agent strictly prefers outcome o_1 to outcome o_2
  - $o_1 \sim o_2$ if the agent is indifferent between outcome o_1 and outcome o_2.
    - $o_1 \sim o_2$ is the same as ($o_1 \succcurlyeq o_2$ and $o_2 \succcurlyeq o_1$)
- These extend to lotteries
  - So, L $\succcurlyeq$ M if the agent weakly prefers lottery L to lottery M

CARDINAL PREFERENCES

- A utility function is a function $u : O \rightarrow \mathbb{R}$.
- It represents a preference relation if

$$o_1 \succcurlyeq o_2 \iff u(o_1) \geq u(o_2)$$

- A utility function expresses not only which of two outcomes is preferred, but also the intensity of preferences: how much better (or worse) one outcome is compared to another.

ORDINAL VS CARDINAL

- Cardinal utilities are very useful (we like working with numbers!)
- But, they appear to be a much stronger assumption than ordinal
- Cardinal utilities assume that levels of absolute satisfaction exist
- von Neumann and Morgenstern proved that under mild conditions, ordinal utilities imply cardinal utilities.

AXIOMS: COMPLETENESS AND TRANSITIVITY

- Axiom 1 (completeness): For every pair of lotteries L and M:

$$L \succ M \text{ or } M \succ L \text{ or } L \sim M$$

- Axiom 2 (transitivity): For any lotteries L_1, L_2, and L_3:

$$\text{If } L_1 \succcurlyeq L_2 \text{ and } L_2 \succcurlyeq L_3 \text{ then } L_1 \succcurlyeq L_3$$

- In other words, the preference relation $\succcurlyeq$ is a weak order

JUSTIFICATION

- No money pumps!
- Suppose that you had preferences such that $o_1 \succ o_2$, $o_2 \succ o_3$ and $o_3 \succ o_1$.
- Starting from o_3, you should be willing to pay (say) a cent to switch to o_2.
- But from o_2 you should be willing to pay a cent to switch to o_1.
- And from o_1 you should be willing to pay a cent to switch back to o_3.
- You can be pumped for money indefinitely, and end up where you began.

AXIOMS: INDEPENDENCE

- Axiom 3 (independence): For any lotteries L_1, L_2, and L_3, and every $\alpha \in (0,1)$:

$L_1 \succcurlyeq L_2$ iff $\alpha L_1 + (1-\alpha)L_3 \succcurlyeq \alpha L_2 + (1-\alpha)L_3$

- Intuition: Mixing both lotteries with the same third lottery should not change which one you prefer.

AXIOMS: CONTINUITY

- Axiom 4(continuity): If $L_1 \succcurlyeq L_2 \succcurlyeq L_3$ for lotteries L_1, L_2, L_3, then there exists an $\alpha \in [0,1]$ such that $L_2 \sim \alpha L_1 + (1-\alpha)L_3$

- Intuition: There is a way to randomize between your worse lottery and your favorite lottery so that you are indifferent between the mix and your medium lottery.

THE VON NEUMANN–MORGENSTERN THEOREM [1947]

- Theorem [von Neumann–Morgenstern]:

A preference relation $\succcurlyeq$ over lotteries satisfies Completeness, Transitivity, Independence and Continuity, if and only if there is a utility function u over outcomes such that, for all lotteries L_1 and L_2:

$$L_1 \succcurlyeq L_2 \iff \mathbb{E}[u(L_1)] \geq \mathbb{E}[u(L_2)]$$

- Here, $\mathbb{E}[u(L)] = \sum_{i=1}^{n} p_i \cdot u(o_i)$ for $\mathrm{L} = [p_1 : o_1, o_2, \ldots, p_n : o_n]$

PROOF STRATEGY

- The proof is constructive.
  - We will build a utility function u (and then check that the expected utility reproduces $\succcurlyeq$).

LEMMA 1: INDEPENDENCE PRESERVES INDIFFERENCE

- Lemma 1. If $L_1 \sim L_2$, then for every lottery L_3 and every $\alpha \in [0,1]$:

$$\alpha L_1 + (1-\alpha)L_3 \sim \alpha L_2 + (1-\alpha)L_3$$

Proof

- The cases $\alpha = 0$ and $\alpha = 1$ are immediate, so assume that $\alpha \in (0,1)$.
- Since $L_1 \sim L_2$, we have both $L_1 \succcurlyeq L_2$ and $L_2 \succcurlyeq L_1$.
- Applying independence to each gives

$$\alpha L_1 + (1-\alpha)L_3 \succcurlyeq \alpha L_2 + (1-\alpha)L_3$$
$$\alpha L_2 + (1-\alpha)L_3 \succcurlyeq \alpha L_1 + (1-\alpha)L_3$$

- Both preferences together are exactly indifference. ■

LEMMA 2: INDEPENDENCE PRESERVES STRICT PREFERENCE

- Lemma 2. If $L_1 \succ L_2$, then for every lottery L_3 and every $\alpha \in (0,1)$:

$$\alpha L_1 + (1-\alpha)L_3 \succ \alpha L_2 + (1-\alpha)L_3$$

Proof

- Suppose, toward a contradiction, that

$$\alpha L_2 + (1-\alpha)L_3 \succcurlyeq \alpha L_1 + (1-\alpha)L_3$$

- Independence runs in both directions, so this would give $L_2 \succcurlyeq L_1$, contradicting $L_1 \succ L_2$.
- Completeness implies that the mixture preference is strict (in the order direction). ■

LEMMA 3: SIMULTANEOUS SUBSTITUTION

- Lemma 3. Suppose $L_i \sim L'_i$ for $i = 1, \ldots, m$, and let $\lambda_i \geq 0$ be weights such that $\sum_i \lambda_i = 1$. Then:

$$\sum_i \lambda_i L_i \sim \sum_i \lambda_i L'_i$$

Proof

- Apply Lemma 1 one component at a time: replace L_1 by L'_1 inside the mixture, then L_2 by L'_2, and so on.
- Every replacement preserves indifference.
- Transitivity then chains the m intermediate steps into one. ■

STEP 1: A BEST AND A WORST OUTCOME

- Because O is finite and $\succcurlyeq$ is complete and transitive, the outcomes can be ranked.
- So, there are outcomes $B, W \in O$ with

$$B \succcurlyeq x \succcurlyeq W \quad for\ every\ x \in O$$

- Call B a best outcome and W a worst outcome
- Assume that $B \succ W$ for now
  - We'll handle $B \sim W$ at the end
- For $\alpha \in [0,1]$ define the reference gamble

$$L_\alpha = \alpha B + (1-\alpha)W$$

- As α increases, we go from the worst to the best outcome.

LEMMA 4: MORE PROBABILITY ON THE BEST OUTCOME IS BETTER

- **Lemma 4.** For all $\alpha, \beta \in [0,1]$:

$$L_\alpha \succcurlyeq L_\beta \iff \alpha \geq \beta$$

$$L_\alpha \succ L_\beta \iff \alpha > \beta$$

Proof

- The endpoints $\alpha = 1$ and $\beta = 0$ are immediate, since L_1 is B and L_0 is W.
- Let $0 < \alpha < 1$
- Apply Lemma 2 to $B \succ W$:

$$L_\alpha = \alpha B + (1-\alpha)W \succ \alpha W + (1-\alpha)W$$

- The RHS is W, so $L_\alpha \succ W$

LEMMA 4: THE COMPARISON BETWEEN TWO GAMBLES

- Take $0 < \beta < \alpha$
- Set $\gamma = 1 - \beta/\alpha \in (0,1)$.
- Since $L_\alpha \succ W$, by Lemma 2 we have

$$\gamma L_\alpha + (1-\gamma)L_\alpha \succ \gamma W + (1-\gamma)L_\alpha$$

- The LHS is just L_α.
- For the RHS:
  - $1 - \gamma = \beta/\alpha$, so substituting $L_\alpha = \alpha B + (1-\alpha)W$ gives

$$\gamma W + (\beta/\alpha)[\alpha B + (1-\alpha)W]$$
$$= \beta B + (1-\beta)W = L_\beta$$

- So $\alpha > \beta$ implies $L_\alpha \succ L_\beta$.
- Reversing the roles covers $\alpha < \beta$, and $\alpha = \beta$ is the same lottery. ■

STEP 2: ASSIGN A UTILITY NUMBER TO EVERY OUTCOME

- Fix an outcome $x \in O$
- Since B is best and W is worst, $B \succcurlyeq x \succcurlyeq W$
- By continuity, there exists an $\alpha \in [0,1]$ such that $x \sim \alpha B + (1-\alpha)W = L_\alpha$
- Define $u(x) = \alpha$
- We must show that α is unique

STEP 2: ASSIGN A UTILITY NUMBER TO EVERY OUTCOME

- Suppose two weights both work: $x \sim L_\alpha$ and $x \sim L_\beta$, for $\alpha \neq \beta$
- Transitivity implies that $L_\alpha \sim L_\beta$
- But, by Lemma 4, if $\alpha > \beta$, $L_\alpha \succ L_\beta$, and if $\alpha < \beta$, $L_\alpha \prec L_\beta$
- Therefore, $\alpha = \beta$ (contradiction)

STEP 2: ASSIGN A UTILITY NUMBER TO EVERY OUTCOME

- Therefore, for every outcome $x \in O$, we have a unique $u(x) \in [0,1]$ such that

$$x \sim u(x)B + \big(1 - u(x)\big)W \qquad (1)$$

- And, in particular, $u(B) = 1$ and $u(W) = 0$

STEP 3: REDUCE AN ARBITRARY LOTTERY TO A BEST–WORST LOTTERY

- Let $L \in \Delta(O)$
- Then $L = \sum_{i=1}^{n} L(o_i)o_i$, by definition
  - for some weights $L(o_i) \geq 0$ such that $\sum_{i=1}^{n} L(o_i) = 1$
  - i.e., for some distribution over outcomes
- By equation (1), for every o_i we have

$$o_i \sim u(o_i)B + \big(1 - u(o_i)\big)W$$

- Using Lemma 3:

$$L = \sum_{i=1}^{n} L(o_i)o_i \sim \sum_{i=1}^{n} L(o_i)\left(u(o_i)B + \big(1 - u(o_i)\big)W\right)$$

STEP 3: REDUCE AN ARBITRARY LOTTERY TO A BEST–WORST LOTTERY

$$\sum_{i=1}^{n} L(o_i)\left(u(o_i)B + \left(1 - u(o_i)\right)W\right)$$

is the same as

$$\left(\sum_{i=1}^{n} L(o_i)u(o_i)\right)B + \left(\sum_{i=1}^{n} L(o_i)\left(1 - u(o_i)\right)\right)W$$

- $\sum_{i=1}^{n} L(o_i) = 1$, so

$$\left(\sum_{i=1}^{n} L(o_i)\left(1 - u(o_i)\right)\right) = 1 - \left(\sum_{i=1}^{n} L(o_i)u(o_i)\right)$$

STEP 3: REDUCE AN ARBITRARY LOTTERY TO A BEST–WORST LOTTERY

- Overall,

$$L \sim \left(\sum_{i=1}^{n} L(o_i)u(o_i)\right)B + \left(1 - \sum_{i=1}^{n} L(o_i)u(o_i)\right)W$$

- Define $U(L) = \sum_{i=1}^{n} L(o_i)u(o_i)$
- Then

$$L \sim U(L)B + \big(1 - U(L)\big)W = L_{U(L)}$$

STEP 4: EXPECTED UTILITY REPRESENTS PREFERENCES

- Let $L, M \in \Delta(O)$ be two arbitrary lotteries
- $L \sim L_{U(L)}$ and $M \sim L_{M(L)}$
- By transitivity $L \succcurlyeq M \Leftrightarrow L_{U(L)} \succcurlyeq L_{U(M)}$
- By Lemma 4, $L_{U(L)} \succcurlyeq L_{U(M)} \Leftrightarrow U(L) \geq U(M)$
- So, $L \succcurlyeq M \Leftrightarrow U(L) \geq U(M)$
- Since $U(Z) = \sum_{i=1}^{n} Z(o_i)u(o_i)$ for a lottery Z, we overall have

$$L \succcurlyeq M \Leftrightarrow \sum_{i=1}^{n} L(o_i)u(o_i) \geq \sum_{i=1}^{n} M(o_i)u(o_i)$$

where $L(o_i)$ and $M(o_i)$ are probability distributions

- i.e., preferences are expressed as expected utilities

CORNER CASE

- Suppose that $B \sim W$
- For every $x \in O$, since B is best and W is worst, $B \succcurlyeq x \succcurlyeq W$
- Since $W \succcurlyeq B$, transitivity gives $x \succcurlyeq W \succcurlyeq B$
- So $x \succcurlyeq B$
- So $x \sim B$
- So $x \sim W$
- For every $x \in O$
- So for every lottery $L \in \Delta(O)$

$$L = \sum_{i=1}^{n} L(o_i)o_i \sim \sum_{i=1}^{n} L(o_i)B = B$$

- Thus, $u(x) = 0$ for every $x \in O$ represents these preferences

WHAT WE'VE SHOWN

- We've shown that, given a preference relation $\succcurlyeq$ over lotteries that satisfies Completeness, Transitivity, Independence and Continuity, there is a utility function u over outcomes such that, for all lotteries L_1 and L_2: $L_1 \succcurlyeq L_2 \iff \mathbb{E}[u(L_1)] \geq \mathbb{E}[u(L_2)]$
- It remains to show the converse
  - If there is a utility function that represents a preference relation, then the preferences satisfy the axioms
  - Homework!

IMPACT

- The von Neumann–Morgenstern theorem is one of the most (if not **the** most) influential theorems in Economics
- Pretty hard to overstate its importance

REALISTIC?

- Are the axioms a realistic description of how people actually choose?
- Allais (1953) suggests not so much
- Scenario 1:
  - L_1 gives 1 million for sure
  - L_2 gives 1 million w.p. 0.89, 5 million w.p. 0.1 and 0 w.p. 0.01
  - Most people choose L_1
- Scenario 2:
  - L_3 gives 1 million w.p. 0.11 and 0 w.p. 0.89
  - L_4 gives 5 million w.p. 0.1 and 0 w.p. 0.9
  - Most people choose L_4
- But $L_1 \succ L_2$ means $0.11\,u(1) > 0.1\,u(5)$, while $L_4 \succ L_3$ means the opposite
- Question (HW?): which axiom is violated??

SUMMARY

- von Neumann–Morgenstern theorem
- Allais paradox



Algorithmic Game Theory

CS 58100 · Game Theory II

Alex Psomas

Fall 2026

GAME THEORY

- “How do rational, self-interested agents act in a given environment?”
- Each agent has a set of possible actions
- Each agent gets rewarded according to the rules of the game
  - Some function that maps vectors of actions to rewards

NORMAL-FORM GAMES

- A set of players $N = \{1, \dots, n\}$
- Each player i has a set of pure strategies (actions) S_i
- $\mathcal{S} = S_1 \times S_2 \times \cdots S_n$
- Each player i has a utility function u_i that maps $\mathcal{S}$ to $\mathbb{R}$
  - $u_i(s_1, \dots, s_n)$ is the utility of agent i when player 1 plays s_1, player 2 plays s_2, and so on

NORMAL-FORM GAMES

- We will mostly think of static games of complete information, where players pick their strategies simultaneously.
  - Static: things do not happen over time
  - Complete information: utility functions, types and risk-aversion (if any) are common knowledge
- Note: perfect information ≠ complete information
  - Perfect: players know what has happened (actions other players took) so far
  - Perfect information makes more sense for sequential games



Reshef Meir • 1st
Associate Prof. at Technion | Game Theory | Social Choice
1w •

Release the prisoners from the Prisoners' Dilemma!

Many of you have heard about the Prisoners' Dilemma, perhaps the most famous example of game theory. So famous, that no one including me can properly remember the contrived story, or explain in a simple way how it relates to the payoff matrix why it even matters. No more.

There is exactly one correct way to present the Dilemma to people, be it colleagues, students, or your Mom:
- You get to choose between getting $100, and some other person that you don't know and will never meet getting $300.
- The other person gets the same choice.

This version is due to Prof. Aumann. That's it. No prisoners. No complex reasoning. No over-thinking what each player should do.
The dominant strategy is immediate. So is the social optimum.
Generalization to n>2 (Tragedy of the Commons) is obvious.

THE SUCKER'S DILEMMA

	Give $300 to other person	Get $100 for yourself
Give $300 to other person	300, 300	0, 400
Get $100 for yourself	400, 0	100, 100

THE SUCKER'S DILEMMA

	Cooperate	Defect
Cooperate	300, 300	0, 400
Defect	400, 0	100, 100

STRATEGIES

- Players can pick a pure strategy
  - Deterministic choice for an action, e.g. ``Defect''
- Players can pick mixed strategies
  - Distribution over actions, e.g. ``Defect With probability ¼, and cooperate with probability ¾''
  - We will overload notation a bit, and use s_i for mixed strategies as well

DOMINANT STRATEGIES

- For player i, strategy s_i dominates strategy s_i' if s_i is better than s_i', no matter what the other players' strategies are
- Weak domination: $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$
- Strong domination: $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}), \forall s_{-i}$
- s_i is a weakly (strictly) dominant strategy for player i if it weakly (strictly) dominates every other strategy

THE SUCKER'S DILEMMA

	Cooperate	Defect
Cooperate	300,300	0 , 400
Defect	400, 0	100, 100

Poll 1

Which of the following is true?

1. Defect is a weakly dominant strategy
2. Defect is a strongly dominant strategy
3. Both 1 & 2
4. Beats me!



PREDICTIONS

- If dominant strategies exist, then perhaps a good prediction of what will be the outcome of a game is that every agent will play a dominant strategy

FAST VS SLOW RUNNER

- Do dominant strategies always exist?
- Fast vs slow runner
  - Compete or stay out?
- Is there a dominant strategy for the slow runner (column player)?

	Compete	Stay out
Compete	2,-2	4,0
Stay out	0,4	0,0

ITERATIVE DELETION

- Maybe there are no dominant strategies
- But, there could be strategies that are always dominated
- Assuming everyone is rational, we can just remove those strategies
  - Maybe the new game has dominant strategies!

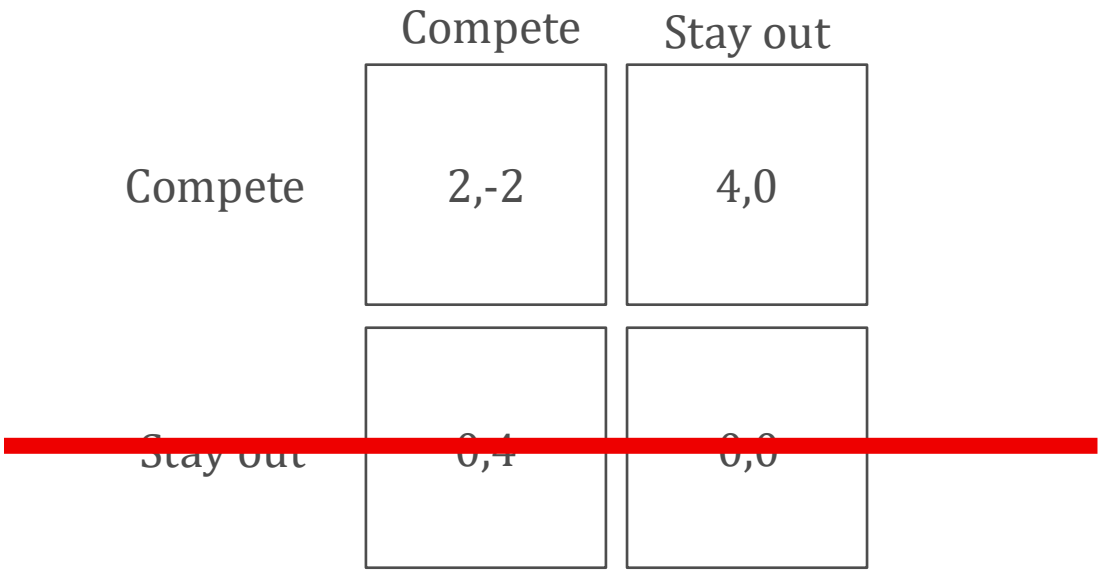
ITERATIVE DELETION

- Fast runner never stays out

	Compete	Stay out
Compete	2,-2	4,0
Stay out	0,4	0,0

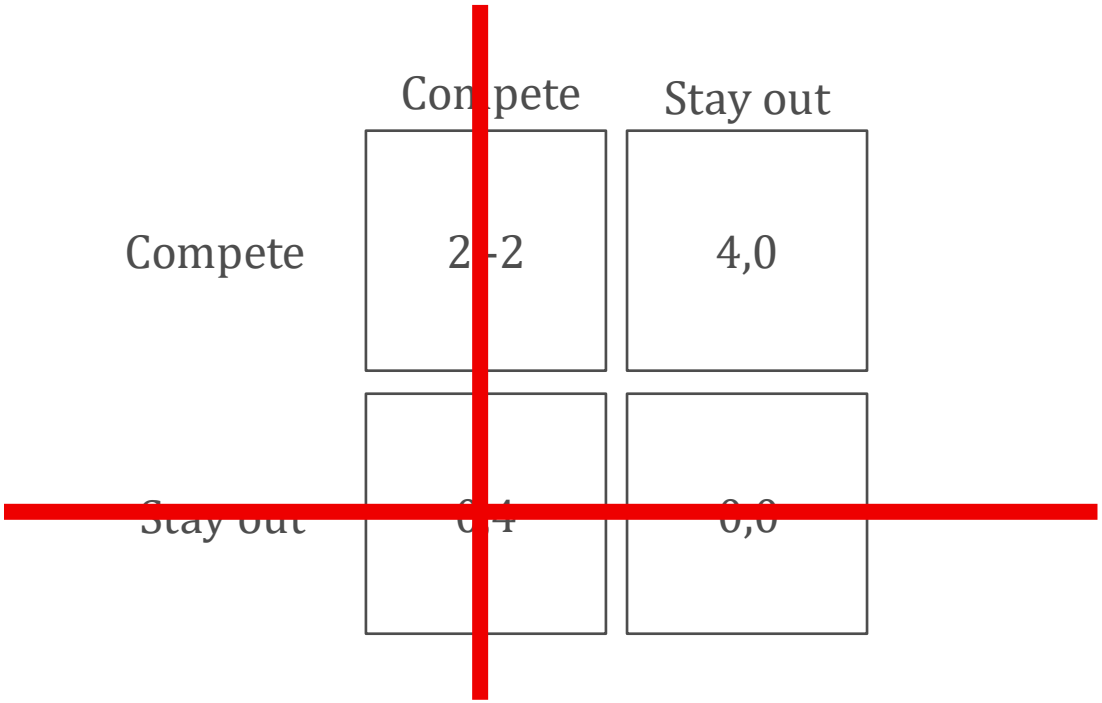
ITERATIVE DELETION

- Fast runner never stays out
- In the new game, slow runner does not want to compete



ITERATIVE DELETION

- Fast runner never stays out
- In the new game, slow runner does not want to compete



ITERATIVE DELETION

- Guess 2/3 of average game:
  - Each student guesses a number between 0 and 100 (inclusive)
  - The student who is closest to 2/3 of the average (of all guesses) wins!

Poll 2

What would you do?



PROFESSOR'S DILEMMA

- If a game has:
  - Dominant strategies, or
  - A unique outcome after iteratively deleting dominated strategies

perhaps we have an accurate prediction of the game's outcome

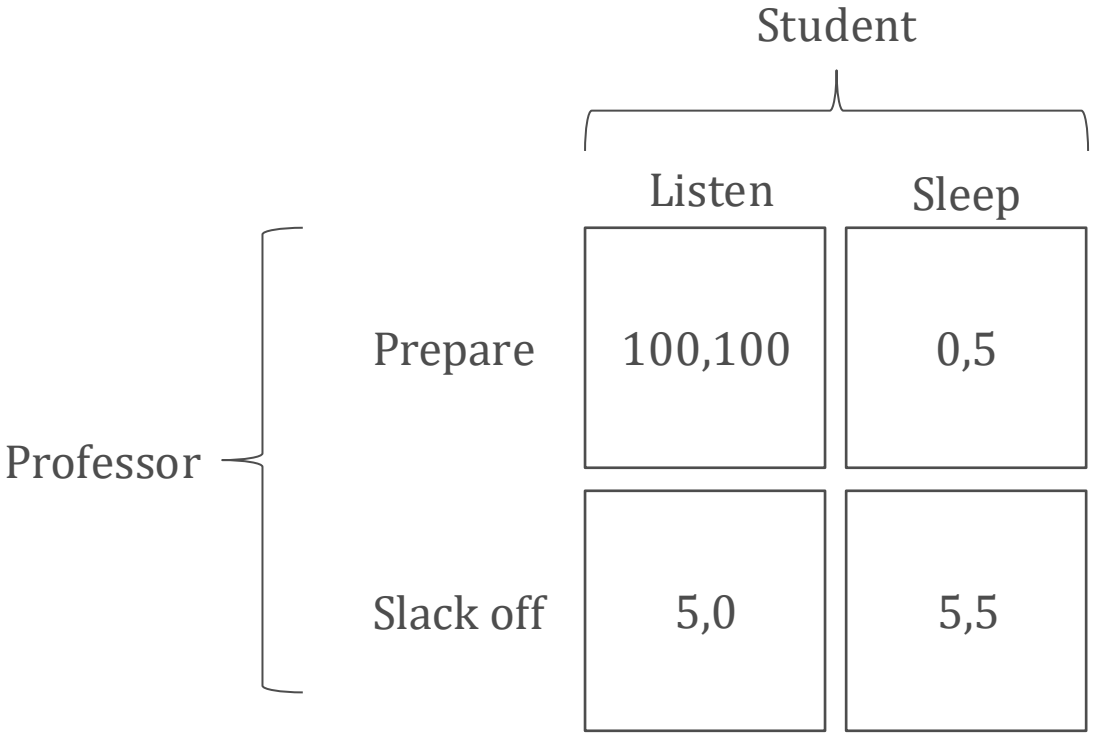
- What if this is not the case?

PROFESSOR'S DILEMMA

- If a game has:
  - Dominant strategies, or
  - A unique outcome after iteratively deleting dominated strategies

perhaps we have an accurate prediction of the game's outcome

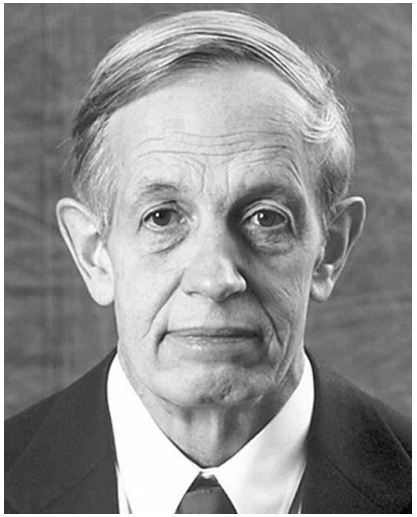
- What if this is not the case?



NASH EQUILIBRIUM

- Distributions $\vec{\sigma} = \sigma_1, \ldots, \sigma_n$ over strategy sets $S_1, \ldots, S_n$ constitute a Nash equilibrium if, for every player i, and every unilateral deviation $s_i' \in S_i$

$$\mathbb{E}_{\vec{s} \sim \vec{\sigma}} [u_i(s_i, s_{-i})] \geq \mathbb{E}_{\vec{s} \sim \vec{\sigma}} [u_i(s_i', s_{-i})]$$



NASH EQUILIBRIUM

- A Nash equilibrium $\vec{\sigma}$ is pure if every σ_i is a pure strategy
  - Each player takes a deterministic action
  - $\sigma_i = s_i$
- Do games always have pure Nash equilibria?

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

- Best response to Rock is Paper
- Best response to Paper is Scissors
- Best response to Scissors is Rock

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

- What's player 2's best response if player 1 plays (0, ½, ½)?
  - That is, P w.p. ½ and S w.p. ½
- Player 1 could play S, and get (expected) utility ½
- Or play P, and get utility -1/2
- Or play R, and get utility 0

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

- What's player 2's best response if player 1 plays $\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$?
- R gives utility 0
- And so does P
- And so does S!
- What is player 2's utility if she also plays $\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$?
- A NE: both play $\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$

MATCHING PENNIES

	Heads	Tails
Heads	+1,-1	-1,+1
Tails	-1,+1	+1,-1

Poll 3

Which are NE in this game?

1. $\left(\left(\frac{1}{2},\frac{1}{2}\right),\left(\frac{1}{2},\frac{1}{2}\right)\right)$
2. $\left(\left(\frac{2}{3},\frac{1}{3}\right),(0,1)\right)$
3. $((1,0),(1,0))$
4. $\left(\left(\frac{1}{2},\frac{1}{2}\right),(1,0)\right)$



INDIFFERENCE PRINCIPLE

- If a mixed strategy σ_i in a Nash equilibrium has support $T_i = \{s_1, \dots, s_k\}$, then every $s_j \in T_i$ gives the same expected payoff to player i (fixing σ_{-i})
- A player is indifferent between all strategies she is randomizing over

STAG HUNT

	Stag	Hare
Stag	4,4	0,2
Hare	2,0	1,1

- Two hunters are playing a game
  - Stag requires both, and feeds them for 4 days each
  - Hare requires 1, but feeds for 2 days

STAG HUNT

	Stag	Hare
Stag	4,4	0,2
Hare	2,0	1,1

- Pure NE:
  - (Stag, Stag)
  - (Hare, Hare)

STAG HUNT

	Stag	Hare
Stag	4,4	0,2
Hare	2,0	1,1

- Mixed NE?
- Stag w.p. p, Hare w.p. $1 - p$
  - Since the game is symmetric let's try a symmetric equilibrium

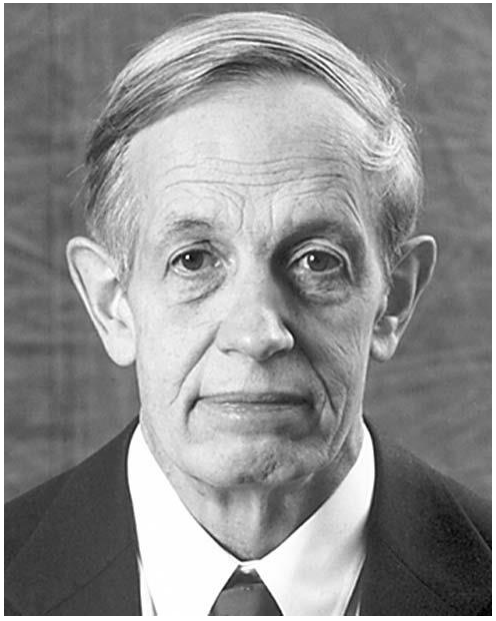
STAG HUNT

	Stag	Hare
Stag	4,4	0,2
Hare	2,0	1,1

- $\mathrm{E}[\text{Stag}] = p \cdot 4 + (1-p) \cdot 0 = 4p$
- $\mathrm{E}[\text{Hare}] = p \cdot 2 + (1-p) \cdot 1 = p + 1$
- $p = 1/3$ makes these two equal
- Claim: $\left(\left(\frac{1}{3}, \frac{2}{3}\right), \left(\frac{1}{3}, \frac{2}{3}\right)\right)$ is a NE

NASH'S THEOREM

- Theorem [Nash, 1950]: In any (finite) game there exists at least one equilibrium.



PROOF OF NASH'S THEOREM

- On the board

SUMMARY

- Definition of games
- Definition of Nash equilibrium
- Existence of NE



Algorithmic Game Theory

CS 58100 · Game Theory III

Alex Psomas

Fall 2026

DOES NE MAKE SENSE?

- Two players
- Strategies are $\{2,3, \dots, 100\}$
- If both choose the same number, this is how much money they get
- If one chooses x, and the other chooses $y < x$, then:
  - The player who picked y gets $y + 2$
  - The player who picked x gets $y - 2$
- What are the equilibria?

DOES NE MAKE SENSE?

- Ultimatum game:
  - Two players: proposer and responder
  - Proposer proposes a way to split a dollar between them $(x, 1-x), x \in [0,1]$
  - Responder accepts or rejects
  - If responder accepts $(x, 1-x)$ the proposer gets x and responder gets $1-x$
  - If responder rejects they both get zero
  - What are the NE?

NASH EQUILIBRIA CRITIQUES

- Rationality is perfect
  - Bounded computation
- Full information about others
  - Bayes-Nash equilibrium
- Humans don’t use expected utility theory
  - Prospect theory
- There are many equilibria, which one will happen in practice?
- Computation (next topic)

COMPUTATIONAL RESULT #1

ZERO SUM GAMES

- Basic definitions
  - Input: $m \times n$ matrix A
  - $A_{i,j}$ is the gain of the row player (and the loss of the column player) when the row player picks pure strategy i and the column player picks pure strategy j
  - Gain of row player for mixed strategies x, y equals $x^T A y = \sum_i \sum_j x_i A_{i,j} y_j$
  - Output: an equilibrium x, y

PURE STRATEGIES

- When the row player picks a pure strategy i, the utility she should expect is $\min_j A_{i,j}$
- So, by picking pure strategies she shouldn't be able to make more than $\max_i \min_j A_{i,j}$
- $\max_i \min_j A_{i,j} \leq \min_j \max_i A_{i,j}$
  - Why?

POLL

$$\max_i \min_j A_{i,j} \leq \min_j \max_i A_{i,j}$$

Poll

Can this inequality be strict?

1. Yes
2. No
3. Beats me!



MIXED STRATEGIES

- When the row player picks a mixed strategy x, the utility she should expect is $\min_y x^T A y$
- So, by picking mixed strategies she shouldn't be able to make more than $\max_x \min_y x^T A\, y$
- Similarly, column player shouldn't lose more than $\min_y \max_x x^T A\, y$

MINMAX THEOREM

- Theorem [Von Neumann]: For every two player zero-sum game there is a value V called the value of the game such that

$$\max_x \min_y x^T A y = V = \min_y \max_x x^T A y$$



COMPUTING AN OPTIMAL STRATEGY

- Example

3	-1
-2	1

- When row picks a mixed strategy $x = (x_1, x_2)$ her utility is

$$z = \min(3x_1 - 2x_2, -x_1 + x_2)$$

  - Should we consider mixed strategies for column?
- When column picks y loss is

$$w = \max(3y_1 - y_2, -2y_1 + y_2)$$

- Solve: Maximize z, with the constraint $x_1 + x_2 \leq 1$!
  - Issues?

TWO OPTIMAL STRATEGIES

- $\max z$
- $z \leq 3x_1 - 2x_2$
- $z \leq -x_1 + x_2$
- $x_1 + x_2 = 1$
- $x_1, x_2 \geq 0$

- $\min w$
- $w \geq 3y_1 - y_2$
- $w \geq -2y_1 + y_2$
- $y_1 + y_2 = 1$
- $y_1, y_2 \geq 0$



DUALS!!!!!!

GENERAL PROOF

- The optimization problem for the row player is
  - $\max z$,
  - $z \leq \sum_i x_i A_{i,j}$, $\forall j$
  - $\sum_i x_i = 1$
- The optimization problem for the column player is its dual!
- By strong duality we get than $z^* = w^*$, or
  - $\max_x \min_j x^T A_{(.,j)} = \min_y \max_i A_{(i,.)} y$
  - $\min_j x^T A_{(.,j)} = \min_y x^T A\, y$
  - $\max_i A_{(i,.)} y = \max_x x\, A\, y$

POSITIVE RESULT #2

ϵ −APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- Input: Two n by n matrices, A and B, for the payoffs of player 1 and player 2.
  - Assume that payoffs are in [0,1]
- Dfn: An ϵ-NE (x, y) satisfies
  - $x^T A y \geq \max_i A_i\, y - \epsilon$, and
  - $y^T B x \geq \max_i B_i x - \epsilon$.
- Goal: Find $\epsilon - NE$ in time $O(n^{\frac{\log n}{\epsilon^2}})$
- Lemma [LMM 03]:
  - There exists an $\epsilon - NE$ where each player uses a strategy that uniformly samples from a multiset (of pure strategies) of size $O(logn/\epsilon^2)$.

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

Proof:

- A NE (x, y) does exist.
- Thought experiment: Sample k times from x and y (with repetition)
  - We view x and y as distributions
- $X =$ multiset of pure strategies after sampling from x
- Y similarly
- $x_i^* = (\#\text{times pure strategy } i \text{ appears in } X)/k$
- $y_i^* = (\#\text{times pure strategy } i \text{ appears in } Y)/k$

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- We want to pick k large enough so that $x^* A y^* \geq \max_i A_i y^* - \epsilon$, and $(y^*)^T B x^* \geq \max_i B_i x^* - \epsilon$
- It suffices if all of the following to hold:
  1. $|A_i y - A_i y^*| \leq \epsilon/3, for\ all\ i$
  2. $|B_j x - B_j x^*| \leq \epsilon/3, for\ all\ j$
  3. $|xAy - x^* A y| \leq \epsilon/3$
  4. $|y^T B x - (y^*)^T B x| \leq \epsilon/3$
- $x^* A y^* \geq x^* A y - \epsilon/3 \geq xAy - 2\epsilon/3 \geq \max_i A_i y - 2\epsilon/3 \geq \max_i A_i y^* - \epsilon$
- Similarly, $(y^*)^T B x^* \geq \max_i B_i x^* - \epsilon$

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- How big should k be?
- Let's first try to bound $\Pr[|A_i y^* - A_i y| \geq \epsilon]$
- $y_j^* = \frac{1}{k}\sum_{\ell=1}^{k} \mathbb{I}\{\, \ell - th\ sample = j \,\}$
- $\mathbb{E}[y_j^*] = y_j$
- $\Pr[|y_j^* - y_j| \geq \epsilon] \leq e^{-2k\epsilon^2}$
- Pick $k = \frac{\log n}{\epsilon^2}$: $\Pr[|y_j^* - y_j| \geq \epsilon] \leq \frac{1}{n^2}$
- So, $\Pr[|A_{ij} y_j^* - A_{ij} y_j| \geq \epsilon] \leq \frac{1}{n^2}$
- By union bound $\Pr[|A_i y^* - A_i y| \geq \epsilon] \leq \frac{1}{n}$
- Similarly for everything else (plus union bound)

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- What have we done so far…
- Sampling $k = O(\frac{\log n}{\epsilon^2})$ times from a NE gives an $\epsilon - NE\ (x^*, y^*)$
- So what?? We don't know the distribution we're supposed to sample from
- Every probability in x^* and y^* is a multiple of $\frac{1}{k}$
- How many such x^* exist?
  - $\binom{n+k-1}{n-1} \approx n^k = n^{\frac{logn}{\epsilon^2}}$
- Try all of them!

COMPUTING A NASH EQUILIBRIUM

Who cares??

If centralized, specially designed algorithms cannot find Nash equilibria, why should we expect distributed, selfish agents to naturally converge to one?

THE PROBLEM

- NASH
  - Input:
    - Number of player n.
    - An enumeration of the strategy set S_p for every player p.
    - The utility function u_p for every player.
    - An approximation requirement ϵ.
  - Output: Compute an ϵ Nash equilibrium
    - Every action that is played with positive probability is an ϵ maximizer (given the other players' strategies)
- Approximation is necessary!
  - There are games with unique irrational equilibria

HOW HARD IS IT TO COMPUTE AN EQUILIBRIUM

- NP-hard perhaps?
- What would a reduction look like?
- Typical reduction: 3SAT to Hamilton cycle
  - Take an instance I of 3SAT
  - Create an instance I' of HC
  - If I' has a Hamiltonian cycle, find a satisfying assignment for I
  - If I' doesn't have Hamiltonian cycle, conclude that there is no satisfying assignment for I

HOW HARD IS IT TO COMPUTE AN EQUILIBRIUM

- 3SAT to NASH?
  - Take an instance I of 3SAT
  - Create an instance I' of NASH
  - If I' has a MNE, find a satisfying assignment for I
  - If I' doesn't have a MNE, conclude that there is no satisfying assignment for I
- All games have a Mixed Nash Equilibrium!

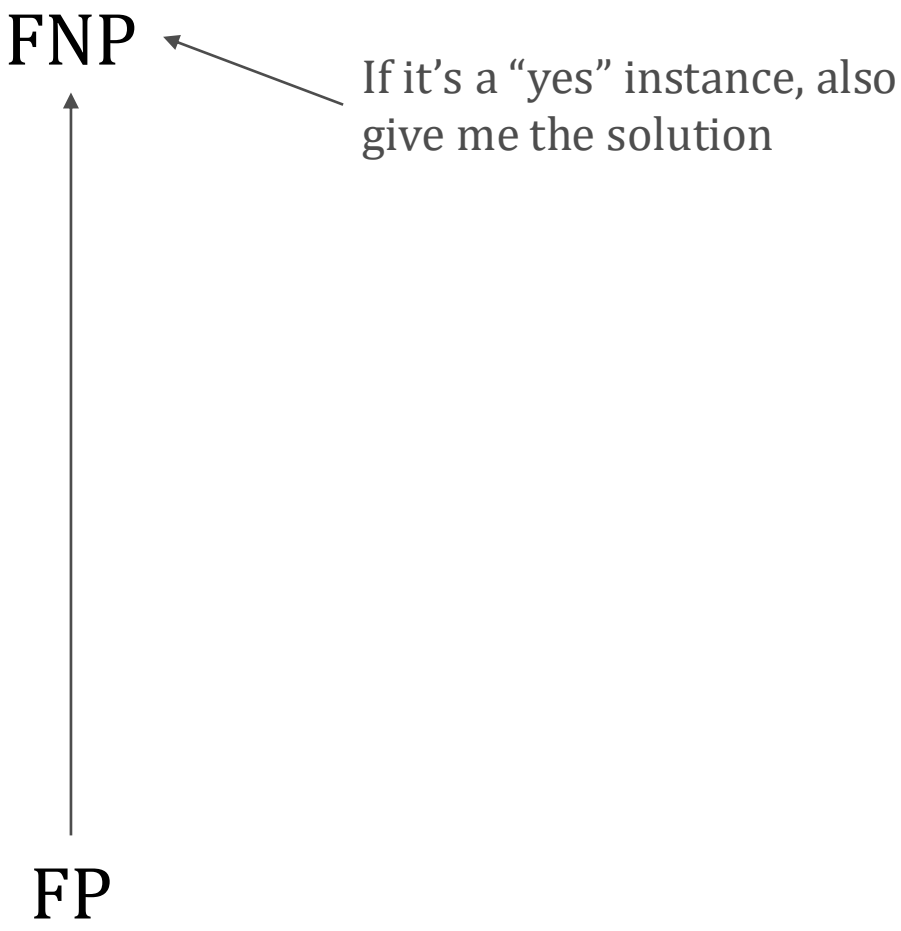
HOW HARD IS IT TO COMPUTE AN EQUILIBRIUM

- What about Pure Nash?
  - Those don't always exist!
  - NP-hard! [Conitzer, Sandholm 2002]
- What about MNE with "social welfare at least x"?
  - NP-hard! [Conitzer, Sandholm 2002]
- What about just MNE?
  - Can't be NP-hard…
  - Doesn't seem to be in P either…
  - Where is it??

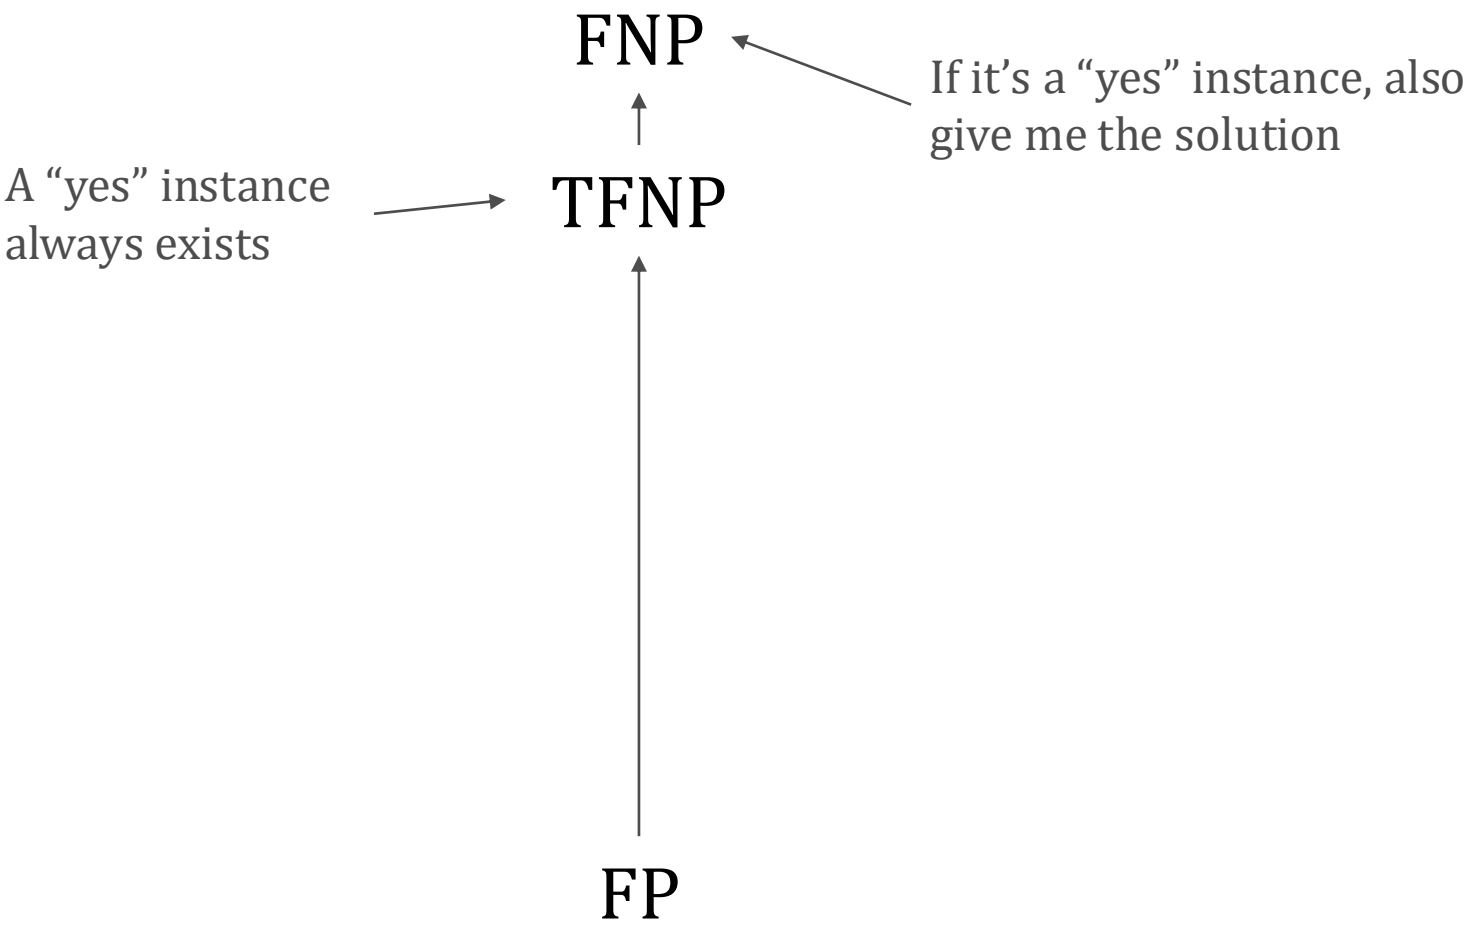
WHICH COMPLEXITY CLASS



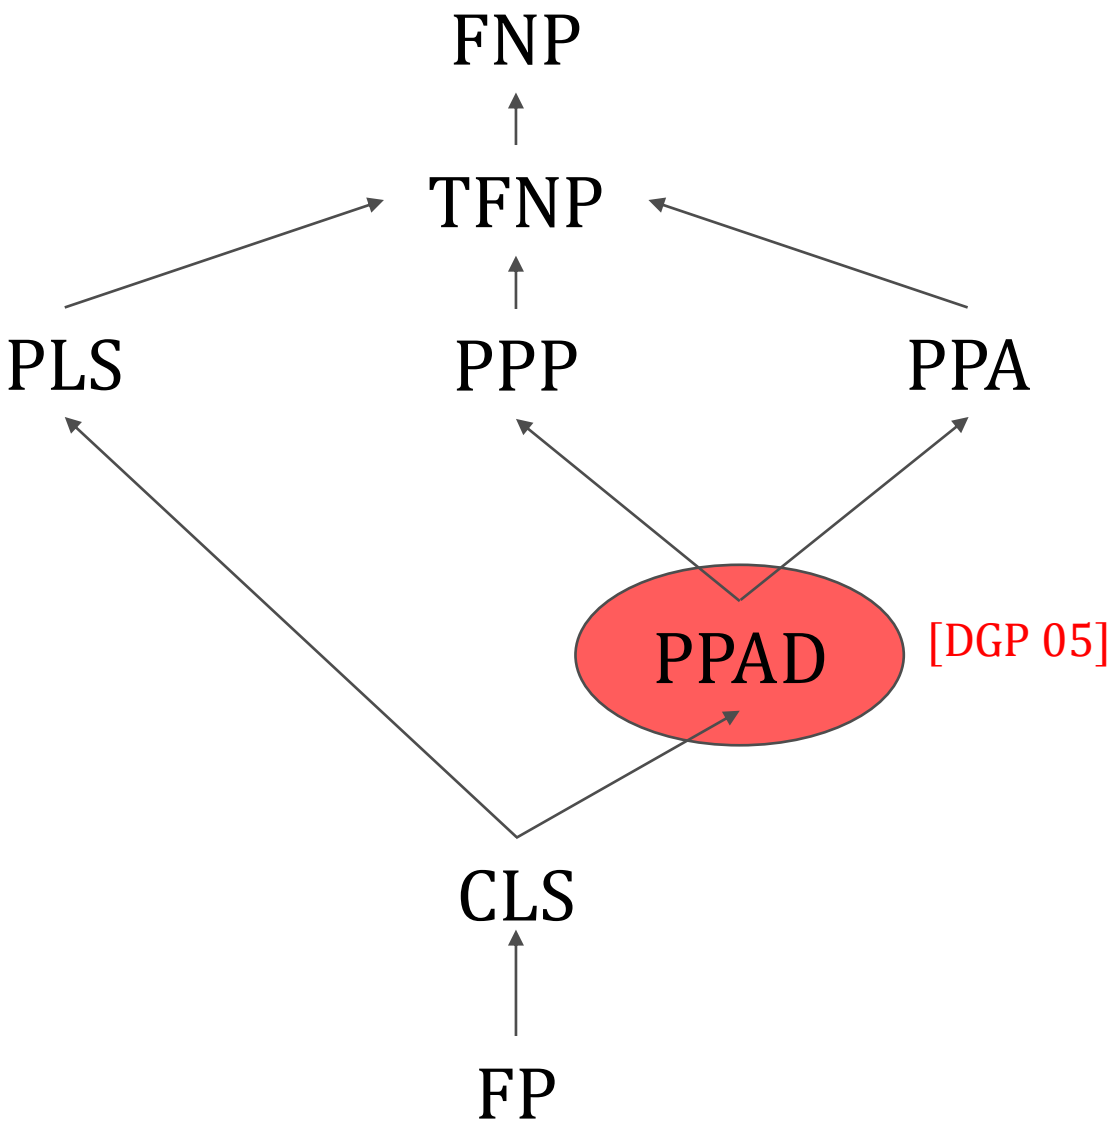
WHICH COMPLEXITY CLASS



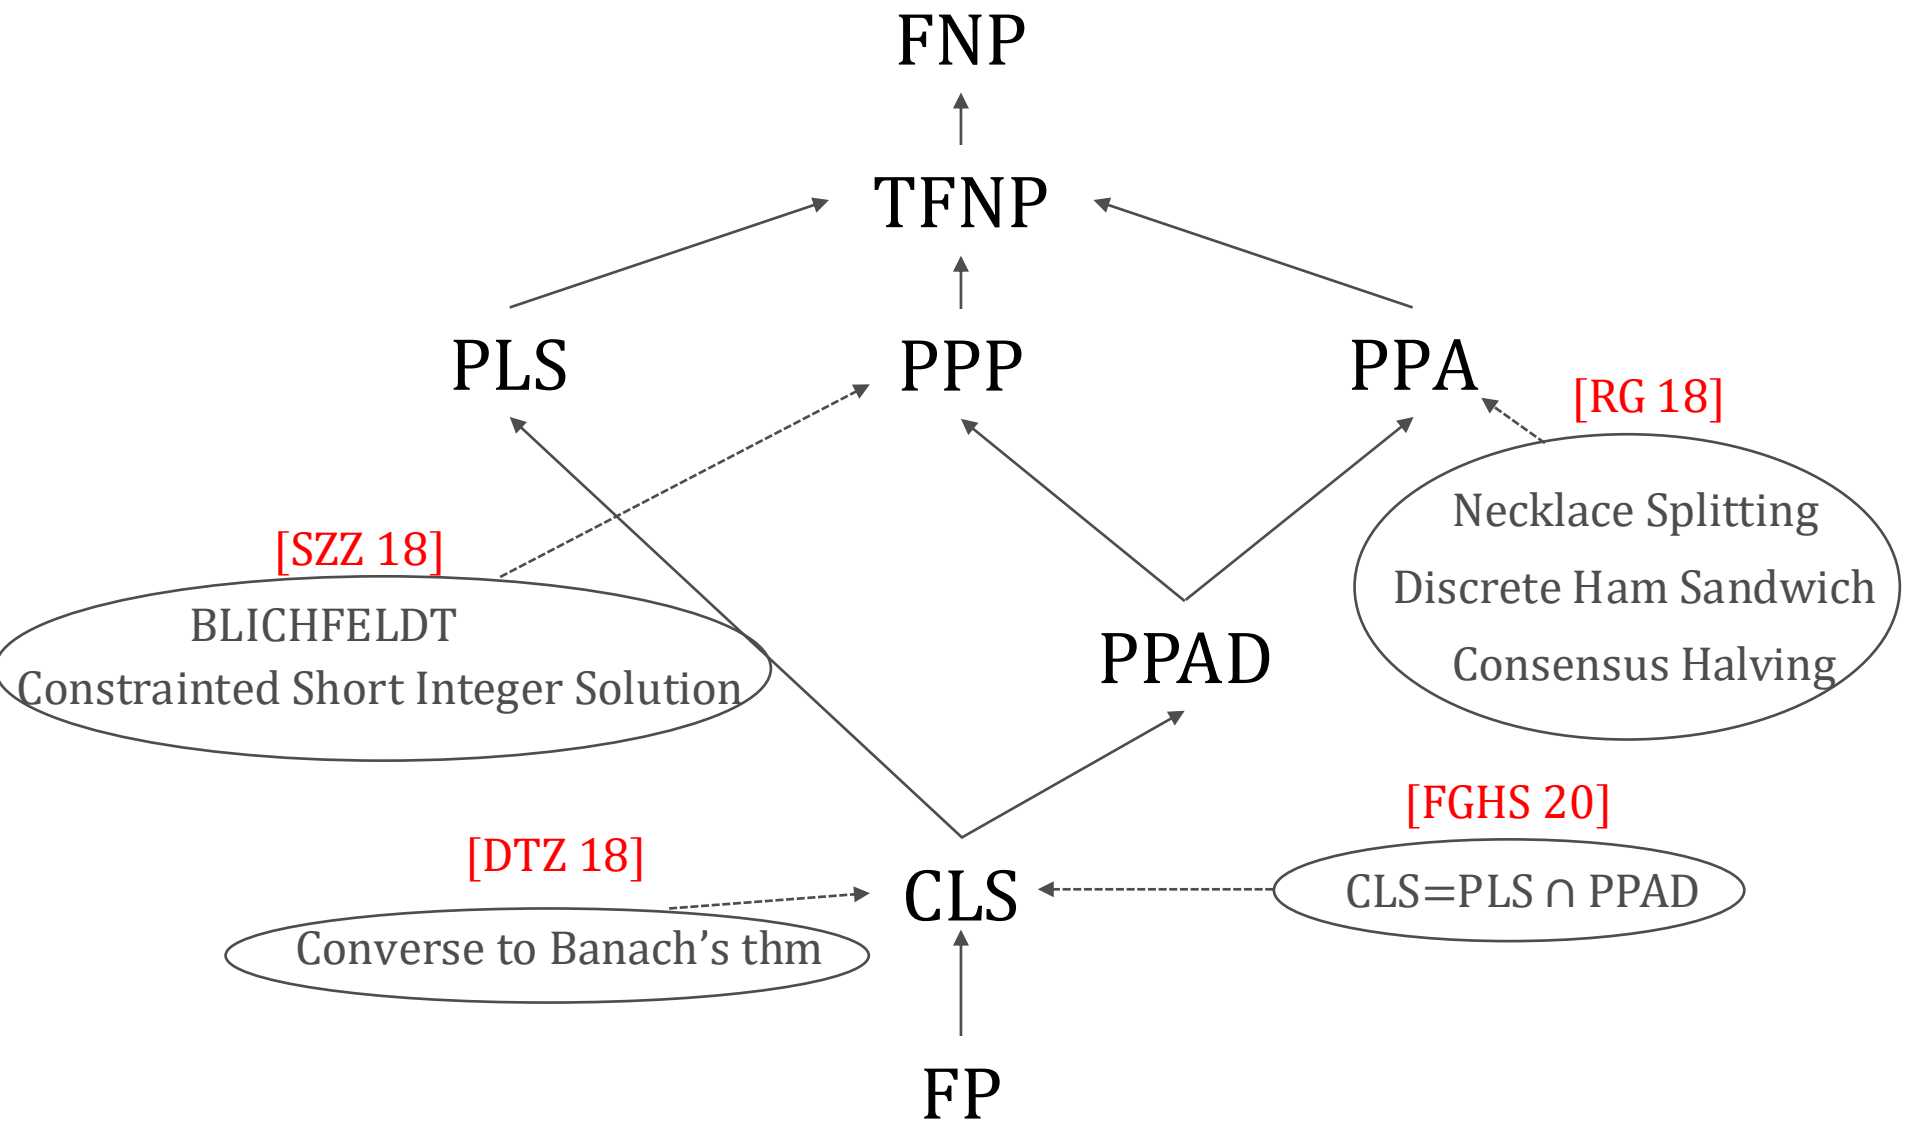
WHICH COMPLEXITY CLASS



WHICH COMPLEXITY CLASS

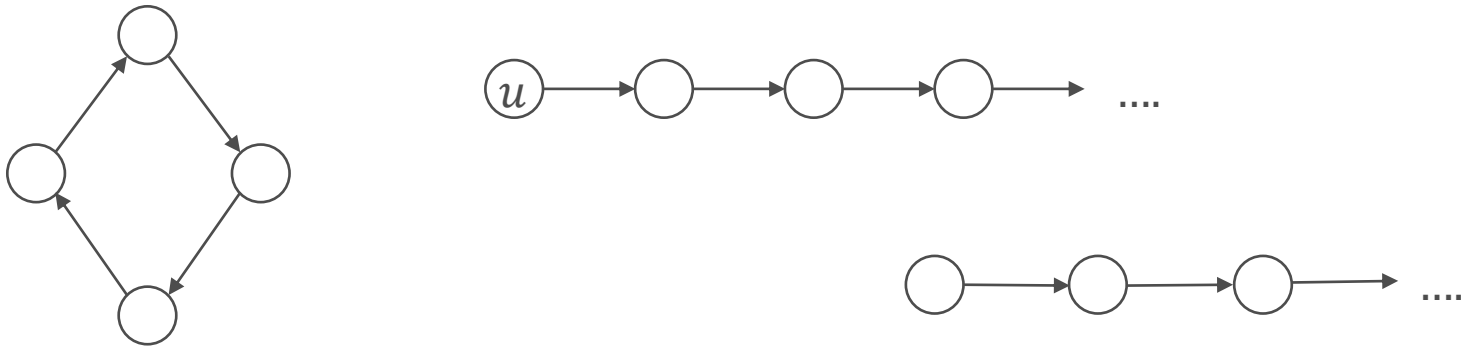


INCIDENTALLY



PPAD

- PPAD: Polynomial Parity Arguments on Directed graphs [Papadimitriou 1994]
- Input: A graph where each vertex has at most in- and out- degree at most 1. A source u.
- Goal: Find a sink or a different source!



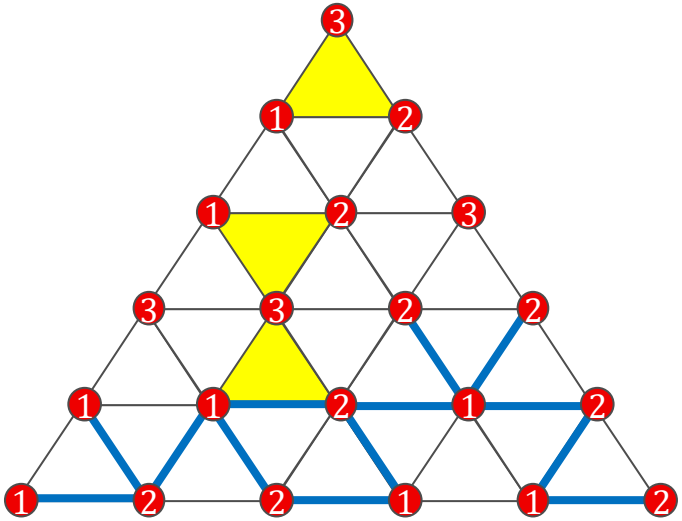
PPAD

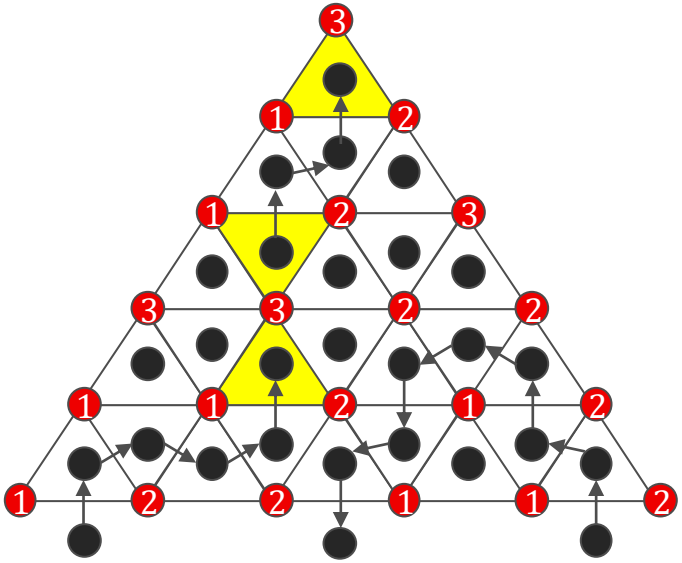
- Why not search the whole graph?
- Graph size is exponential!
- **EndOfALine**: Given two circuits S and P, with m input bits and m output bits each, such that $P(0^m) = 0^m \neq S(0^m)$, find an input $x \in \{0,1\}^m$ such that $P\big(S(x)\big) \neq$ x or $S\big(P($ 0^m.
- PPAD the set of problems reducible to EndOfALine.

WHAT DOES MNE HAVE TO DO WITH ALL THIS?

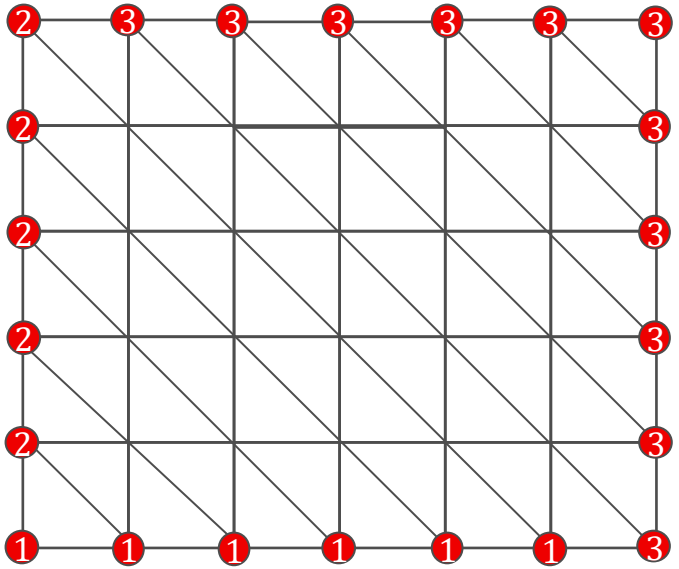
- Nash's proof that every finite game has a MNE uses a fixed point theorem argument, **Brouwer's** fixed point theorem.
- The proof of Brouwer's fixed point theorem uses **Sperner's** Lemma.
- The proof of Sperner's Lemma is at its heart an exponential time **path-following** algorithm!

SPERNER'S LEMMA



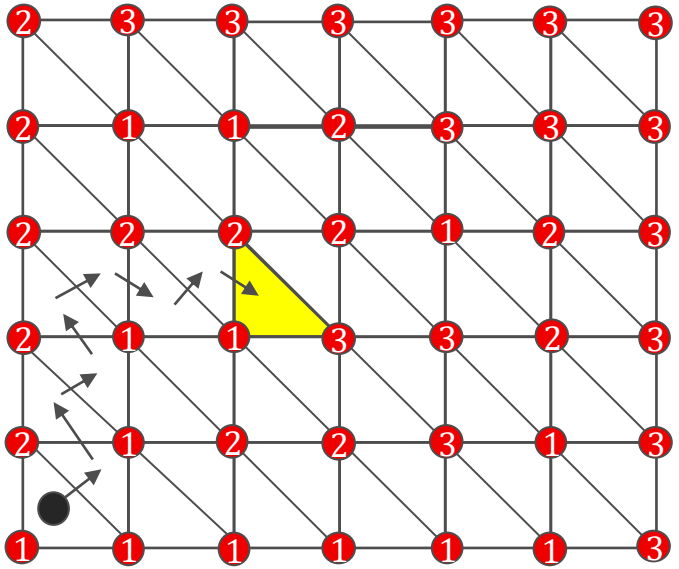


SPERNER'S LEMMA



- 2D Sperner:
  - Input: The description of a poly-time Turing machine f that gives a valid coloring. $f(p) \in \{ 0, 1, 2 \}$, where p is a node.
  - Output: A trichromatic triangle
- 2D-Sperner $\in$ PPAD
  - Obvious reduction.
- 2D-Sperner is PPAD-complete [CD 2006]

SPERNER'S LEMMA

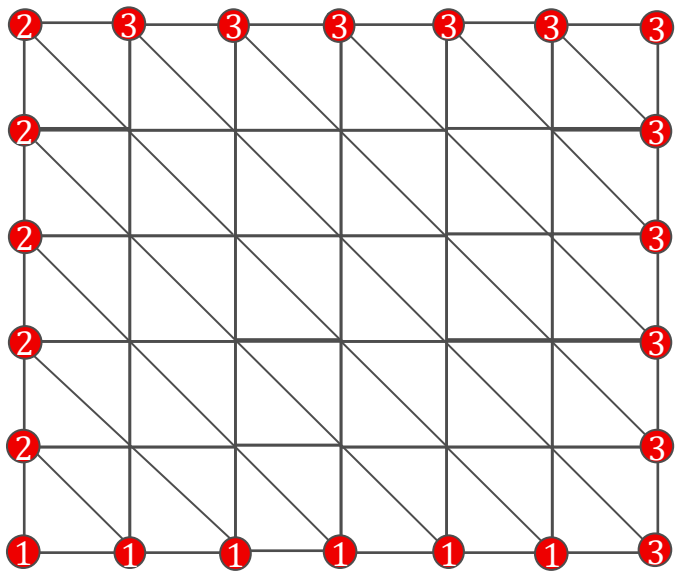


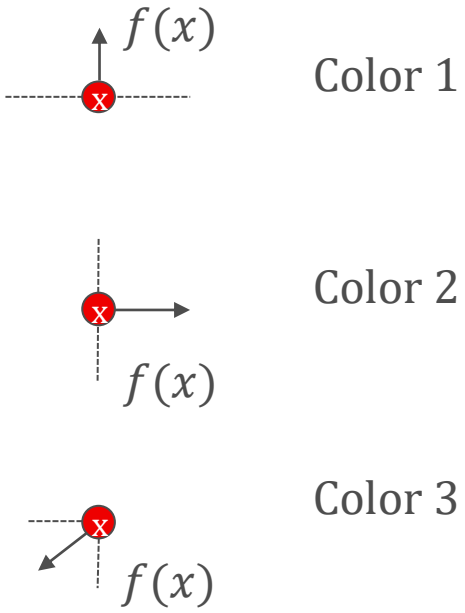
- 2D Sperner:
  - Input: The description of a poly-time Turing machine f that gives a valid coloring. $f(p) \in \{ 0, 1, 2 \}$, where p is a node.
  - Output: A trichromatic triangle
- 2D-Sperner $\in$ PPAD
  - Obvious reduction.
- 2D-Sperner is PPAD-complete [CD 2006]

BROUWER'S FIXED POINT THEOREM

- Thm: Every continuous function f from a closed, convex and compact set C to itself has a fixed point, i.e. a point x_0 such that $f(x_0) = x_0$
- Proof (for $C = [0,1]^2$)
  - Subdivide C into tiny triangles
  - Color the edges like before.
  - For the internal nodes $x = (x_1, x_2)$:
    - If $\big(f(x)\big)_2 = f_2(x) \geq x_2$, color x with color 1
    - If $f_1(x) \geq x_1$, color x with color 2
    - If $f_1(x) \leq x_1$ and $f_2(x) \leq x_2$, color x with color 3
    - If more than 1 condition is met, pick an arbitrary color

BROUWER'S FIXED POINT THEOREM



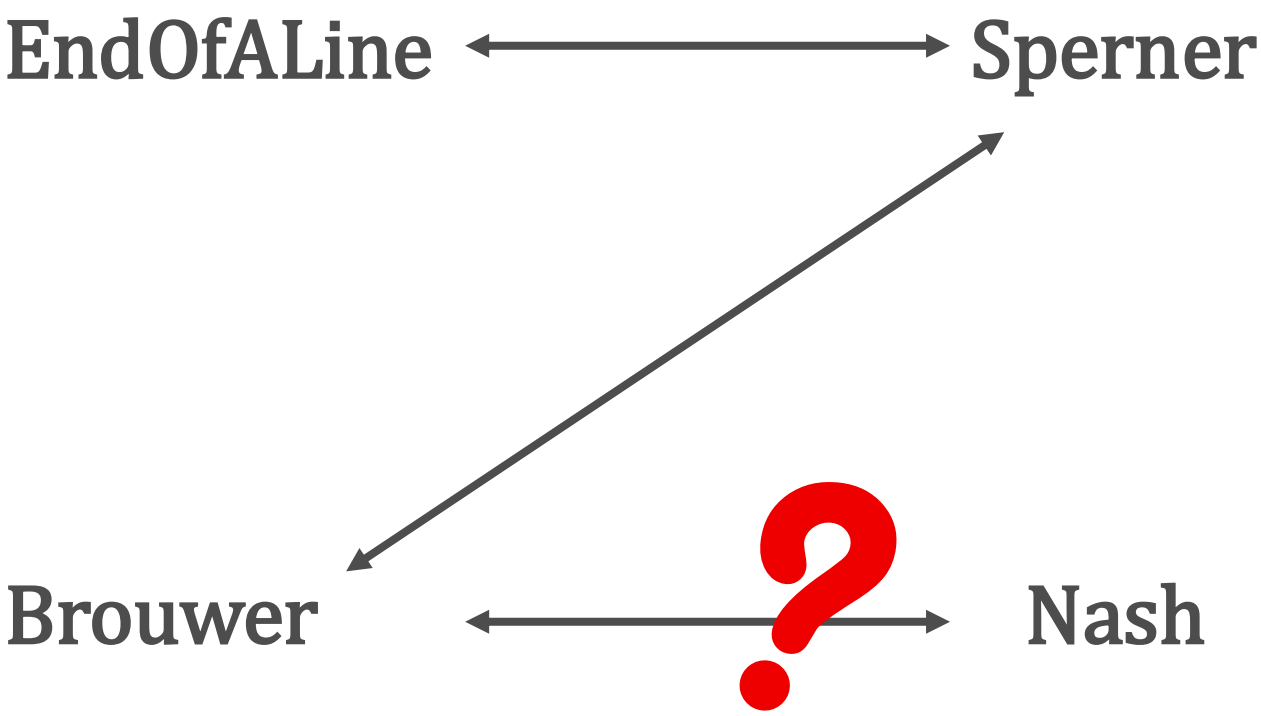


- Color 1 = $f(x)$ farther from bottom than x
- Color 2 = $f(x)$ farther from left side than x
- Color 3 = $f(x)$ farther from top and right side than x
- Trichromatic triangle (in the limit) = $f(x)$ farther from all sides than $x = x$ is a fixed point!

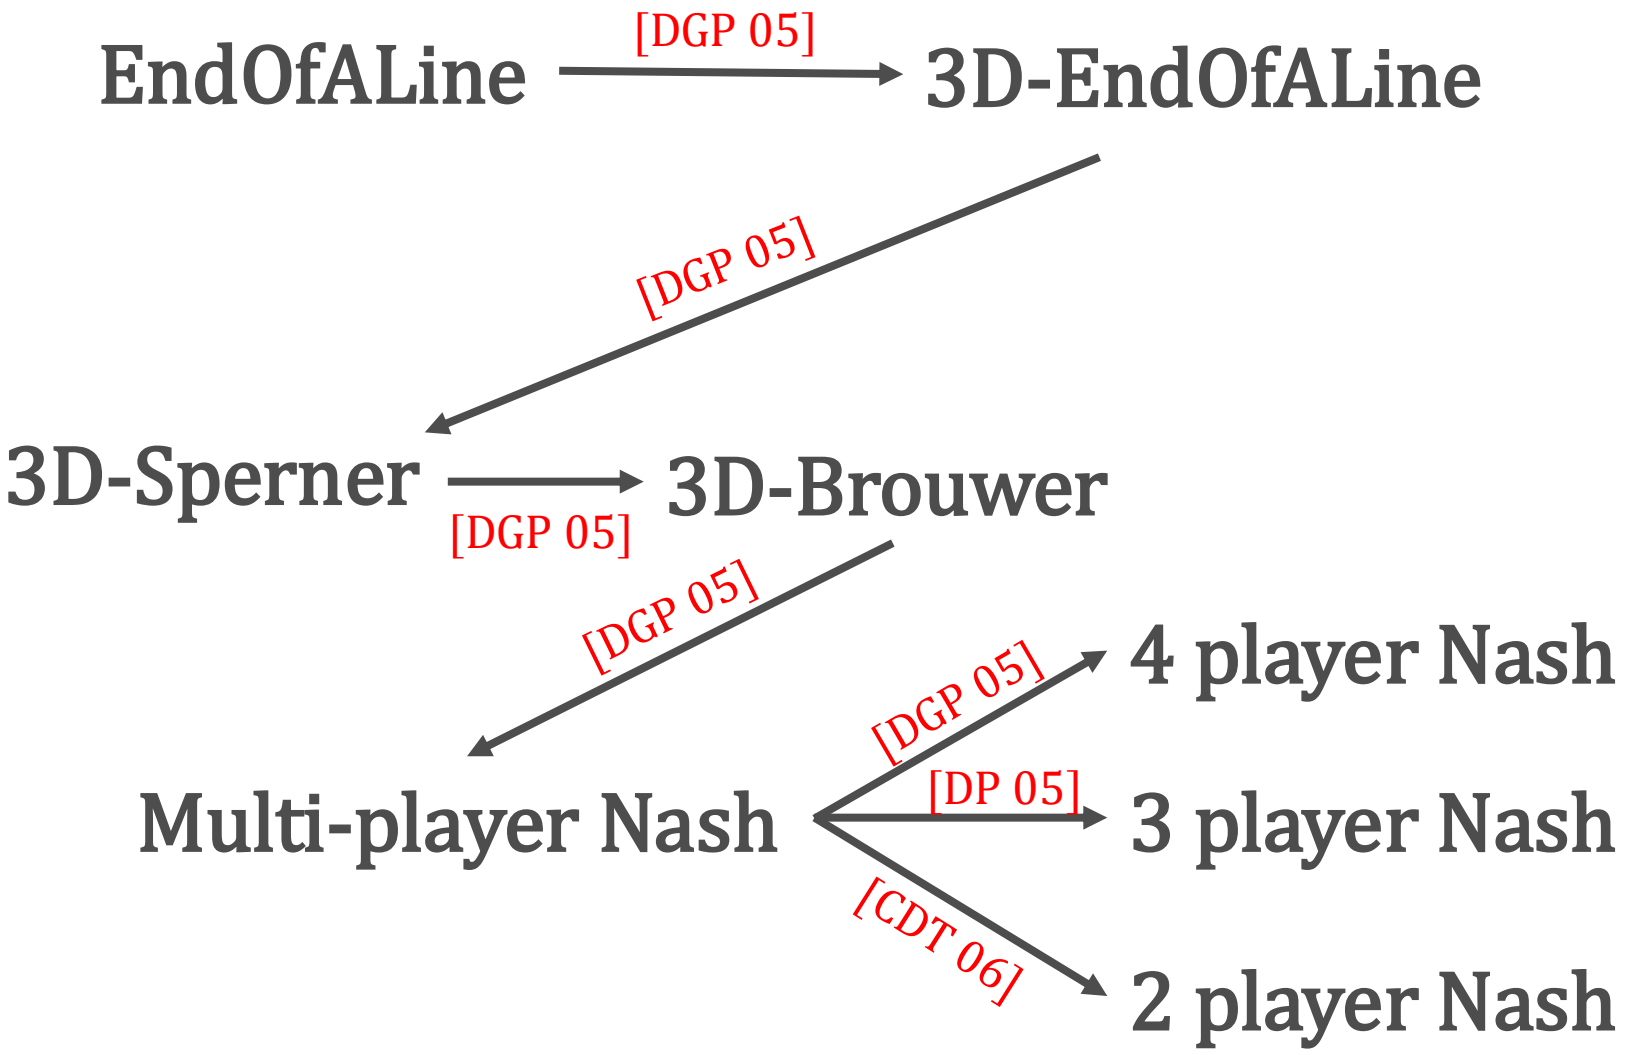
BROUWER'S FIXED POINT THEOREM

- The fixed point could be irrational!
  - We need approximation!
- Brouwer computational problem
  - Input: An algorithm A that evaluates a continuous function f from $[0,1]^n$ to $[0,1]^n$. An approximation ϵ. A Lipschitz constant c that f is claimed to satisfy.
  - Output: x such that $|f(x) - x| < \epsilon$, or a violation of the assumptions
    - $A(x)$ outside $[0,1]^n$, or $|f(x) - f(y)| > c|x - y|$
- Brouwer is PPAD-complete [DGP 05]

STORY SO FAR



THE ACTUAL STORY



THE WHOLE STORY

- Exponential approximation is PPAD complete for 3 players [DGP 06]
- Polynomial approximation is PPAD complete for 2 player NASH [CDT 06]
- Constant approximation is PPAD complete for n players [Rubinstein 15]
- Quasi-polynomial time algorithm for ϵ approximation for 2 player [LMM 03]
- Assuming ETH for PPAD, ϵ approximation takes time $2^{\Omega(n)}$ [Rubinstein 16]

REFERENCES

- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. 2009. The complexity of computing a Nash equilibrium. Commun. ACM
- Chen, X., Deng, X., and Teng, S.-H. 2009. Settling the complexity of computing two-player Nash equilibria. J. ACM
- Rubinstein, A. Inapproximability of Nash equilibrium. STOC 2015
- Rubinstein, A. Settling the Complexity of Computing Approximate Two-Player Nash Equilibria. FOCS 2016
- Lipton, R. J., Markakis, E., and Mehta, A. Playing large games using simple strategies. EC 2003



Algorithmic Game Theory

CS 58100 · Game Theory IV

Alex Psomas

Fall 2026

PROJECTS

- I posted a doc with projects
- Let's do a check-in by next Fri.
- Send Abhiram and I and email with your proposed project/survey
  - Surveys that might need two people
    - Dynamics, chain-recurrent equilibria
    - Complexity of Nash and PPAD

POSITIVE RESULT #2

ϵ −APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- Input: Two n by n matrices, A and B, for the payoffs of player 1 and player 2.
  - Assume that payoffs are in [0,1]
- Dfn: An ϵ-NE (x, y) satisfies
  - $x^T A y \geq \max_i A_i\, y - \epsilon$, and
  - $y^T B x \geq \max_i B_i x - \epsilon$.
- Goal: Find $\epsilon - NE$ in time $O(n^{\frac{\log n}{\epsilon^2}})$
- Lemma [LMM 03]:
  - There exists an $\epsilon - NE$ where each player uses a strategy that uniformly samples from a multiset (of pure strategies) of size $O(logn/\epsilon^2)$.

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

Proof:

- A NE (x, y) does exist.
- Thought experiment: Sample k times from x and y (with repetition)
  - We view x and y as distributions
- $X =$ multiset of pure strategies after sampling from x
- Y similarly
- $x_i^* = (\#\text{times pure strategy } i \text{ appears in } X)/k$
- $y_i^* = (\#\text{times pure strategy } i \text{ appears in } Y)/k$

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- We want to pick k large enough so that $x^* A y^* \geq \max_i A_i y^* - \epsilon$, and $(y^*)^T B x^* \geq \max_i B_i x^* - \epsilon$
- It suffices if all of the following to hold:
  1. $|A_i y - A_i y^*| \leq \epsilon/3, for\ all\ i$
  2. $|B_j x - B_j x^*| \leq \epsilon/3, for\ all\ j$
  3. $|xAy - x^* A y| \leq \epsilon/3$
  4. $|y^T B x - (y^*)^T B x| \leq \epsilon/3$
- $x^* A y^* \geq x^* A y - \epsilon/3 \geq xAy - 2\epsilon/3 \geq \max_i A_i y - 2\epsilon/3 \geq \max_i A_i y^* - \epsilon$
- Similarly, $(y^*)^T B x^* \geq \max_i B_i x^* - \epsilon$

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- How big should k be?
- Let's first try to bound $\Pr[|A_i y^* - A_i y| \geq \epsilon]$
- $y_j^* = \frac{1}{k}\sum_{\ell=1}^{k} \mathbb{I}\{\, \ell - th\ sample = j \,\}$
- $\mathbb{E}[y_j^*] = y_j$
- $\Pr[|y_j^* - y_j| \geq \epsilon] \leq e^{-2k\epsilon^2}$
- Pick $k = \frac{\log n}{\epsilon^2}$: $\Pr[|y_j^* - y_j| \geq \epsilon] \leq \frac{1}{n^2}$
- So, $\Pr[|A_{ij} y_j^* - A_{ij} y_j| \geq \epsilon] \leq \frac{1}{n^2}$
- By union bound $\Pr[|A_i y^* - A_i y| \geq \epsilon] \leq \frac{1}{n}$
- Similarly for everything else (plus union bound)

ϵ APPROXIMATION: NON-ZERO SUM TWO PLAYER GAMES

- What have we done so far…
- Sampling $k = O(\frac{\log n}{\epsilon^2})$ times from a NE gives an $\epsilon - NE$ (x^*, y^*)
- So what?? We don't know the distribution we're supposed to sample from
- Every probability in x^* and y^* is a multiple of $\frac{1}{k}$
- How many such x^* exist?
  - $\binom{n+k-1}{n-1} \approx n^k = n^{\frac{logn}{\epsilon^2}}$
- Try all of them!

COMPUTING A NASH EQUILIBRIUM

Who cares??

If centralized, specially designed algorithms cannot find Nash equilibria, why should we expect distributed, selfish agents to naturally converge to one?

THE PROBLEM

- NASH
  - Input:
    - Number of player n.
    - An enumeration of the strategy set S_p for every player p.
    - The utility function u_p for every player.
    - An approximation requirement ϵ.
  - Output: Compute an ϵ Nash equilibrium
    - Every action that is played with positive probability is an ϵ maximizer (given the other players' strategies)
- Approximation is necessary!
  - There are games with unique irrational equilibria

HOW HARD IS IT TO COMPUTE AN EQUILIBRIUM

- NP-hard perhaps?
- What would a reduction look like?
- Typical reduction: 3SAT to Hamilton cycle
  - Take an instance I of 3SAT
  - Create an instance I' of HC
  - If I' has a Hamiltonian cycle, find a satisfying assignment for I
  - If I' doesn't have Hamiltonian cycle, conclude that there is no satisfying assignment for I

HOW HARD IS IT TO COMPUTE AN EQUILIBRIUM

- 3SAT to NASH?
  - Take an instance I of 3SAT
  - Create an instance I' of NASH
  - If I' has a MNE, find a satisfying assignment for I
  - If I' doesn't have a MNE, conclude that there is no satisfying assignment for I
- All games have a Mixed Nash Equilibrium!

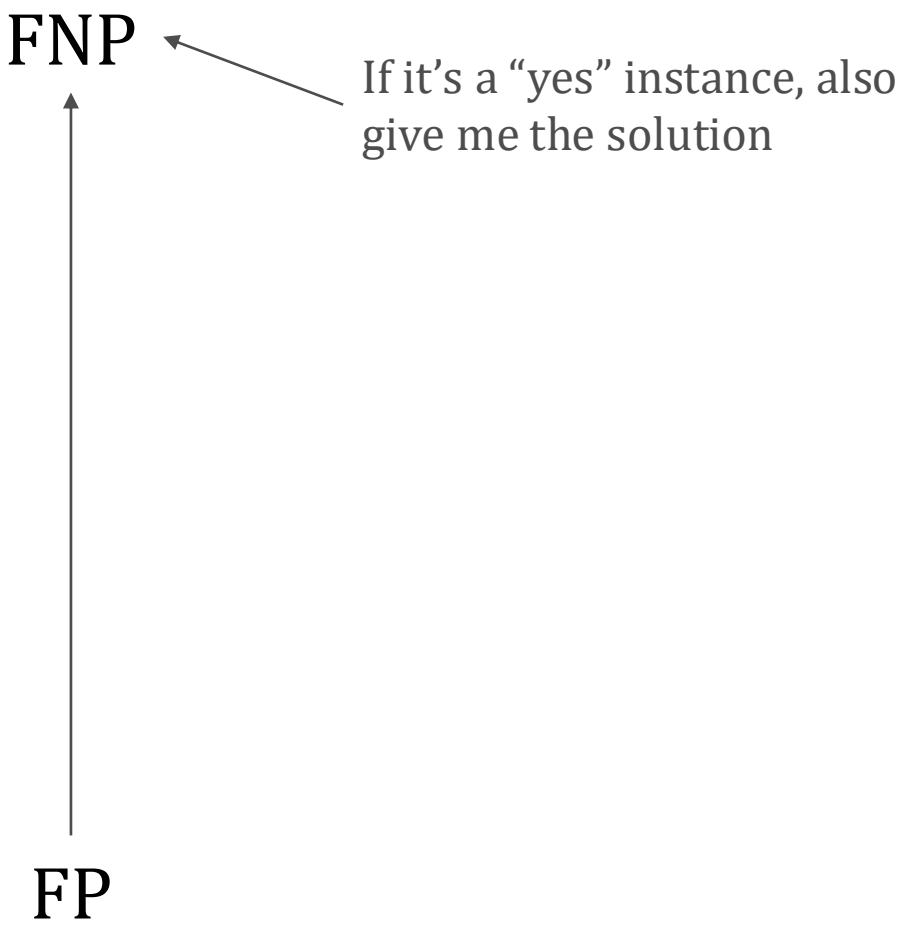
HOW HARD IS IT TO COMPUTE AN EQUILIBRIUM

- What about Pure Nash?
  - Those don't always exist!
  - NP-hard! [Conitzer, Sandholm 2002]
- What about MNE with "social welfare at least x"?
  - NP-hard! [Conitzer, Sandholm 2002]
- What about just MNE?
  - Can't be NP-hard…
  - Doesn't seem to be in P either…
  - Where is it??

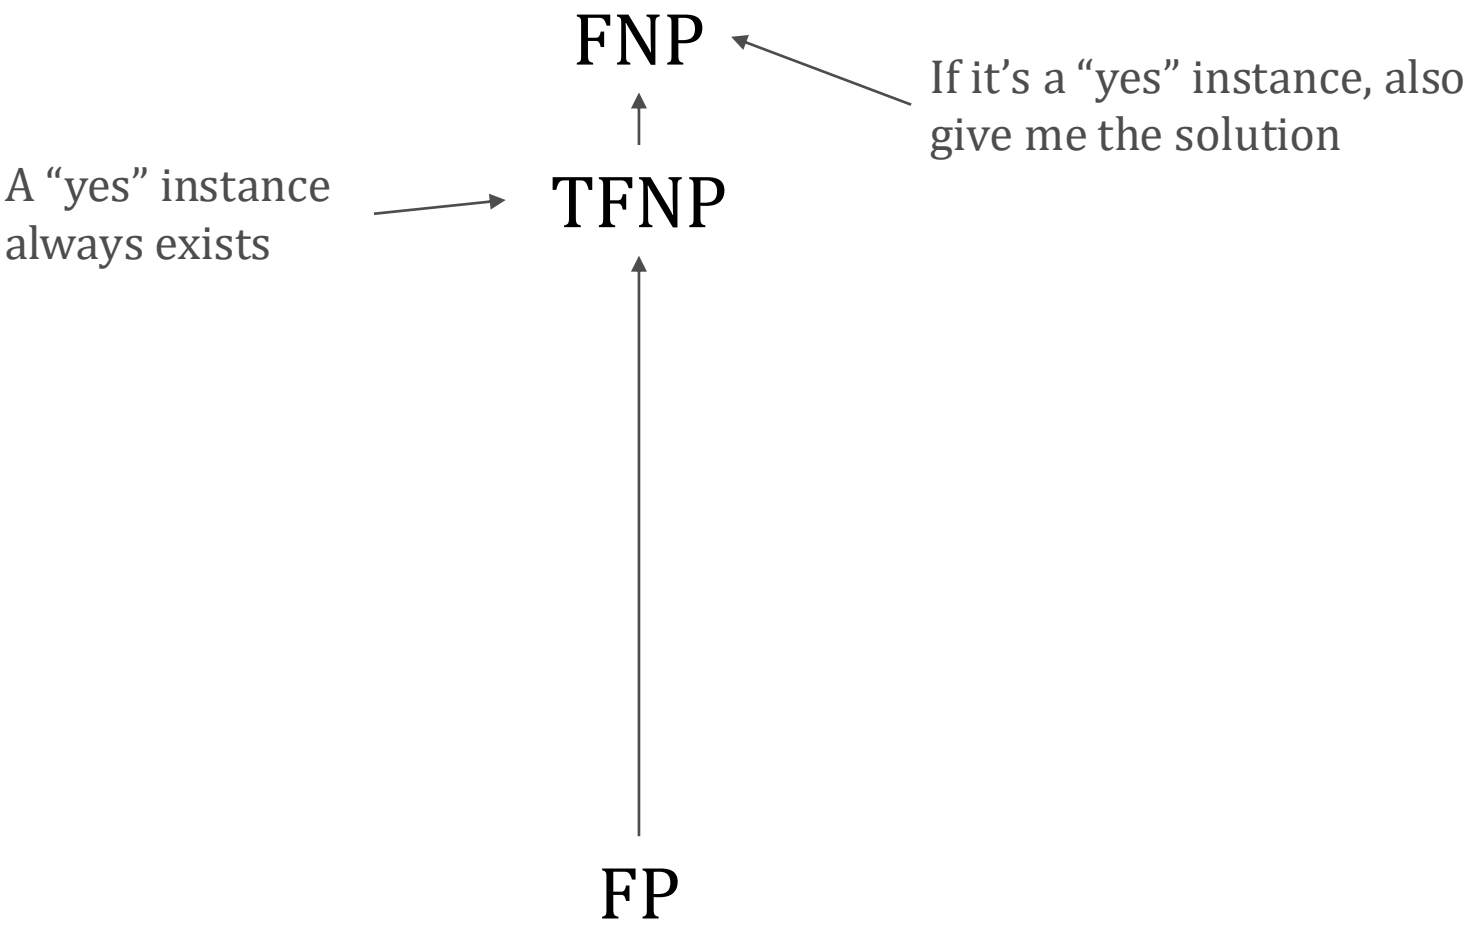
WHICH COMPLEXITY CLASS



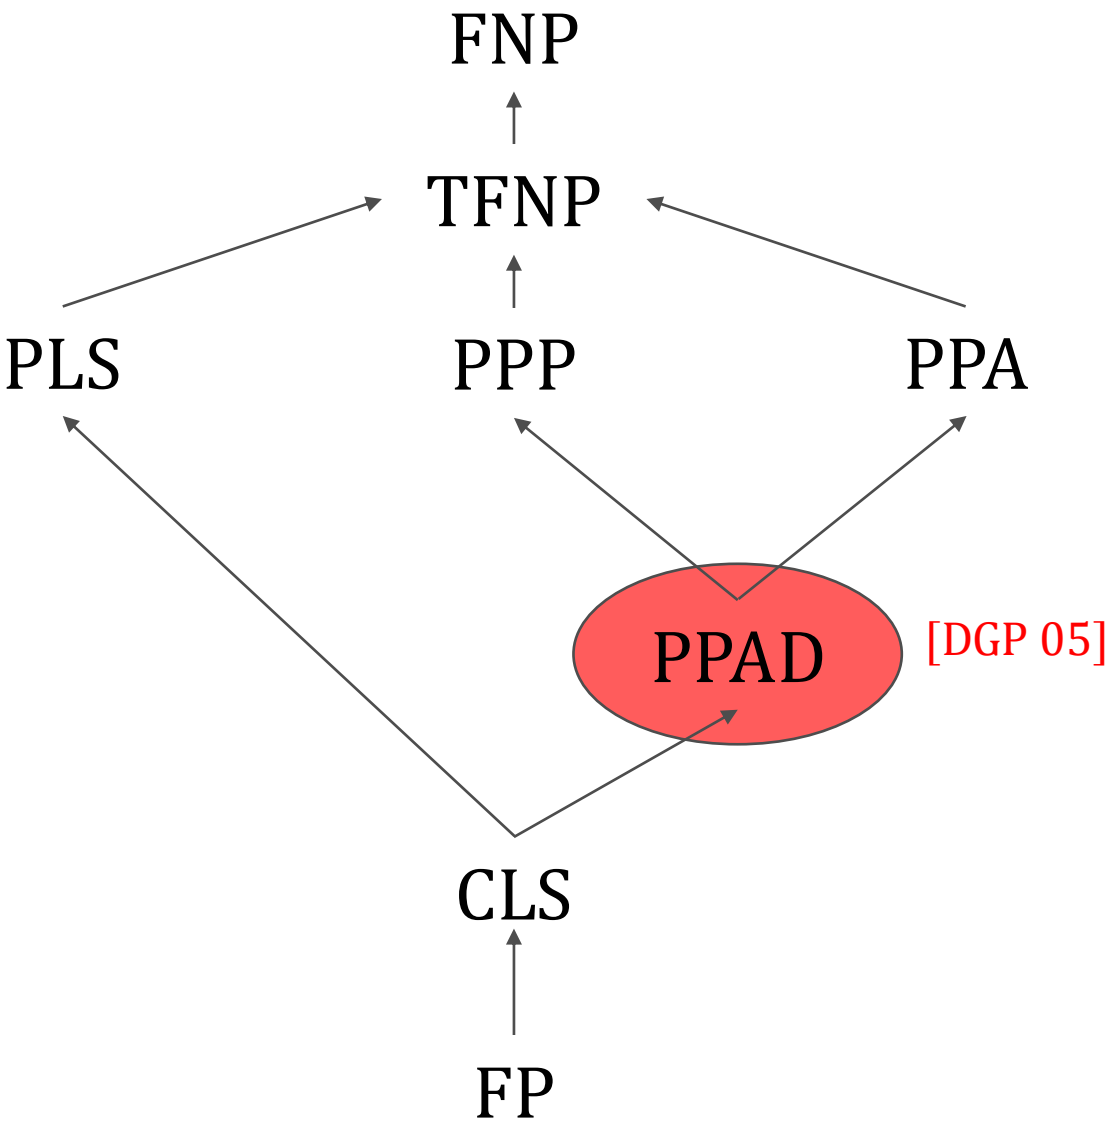
WHICH COMPLEXITY CLASS



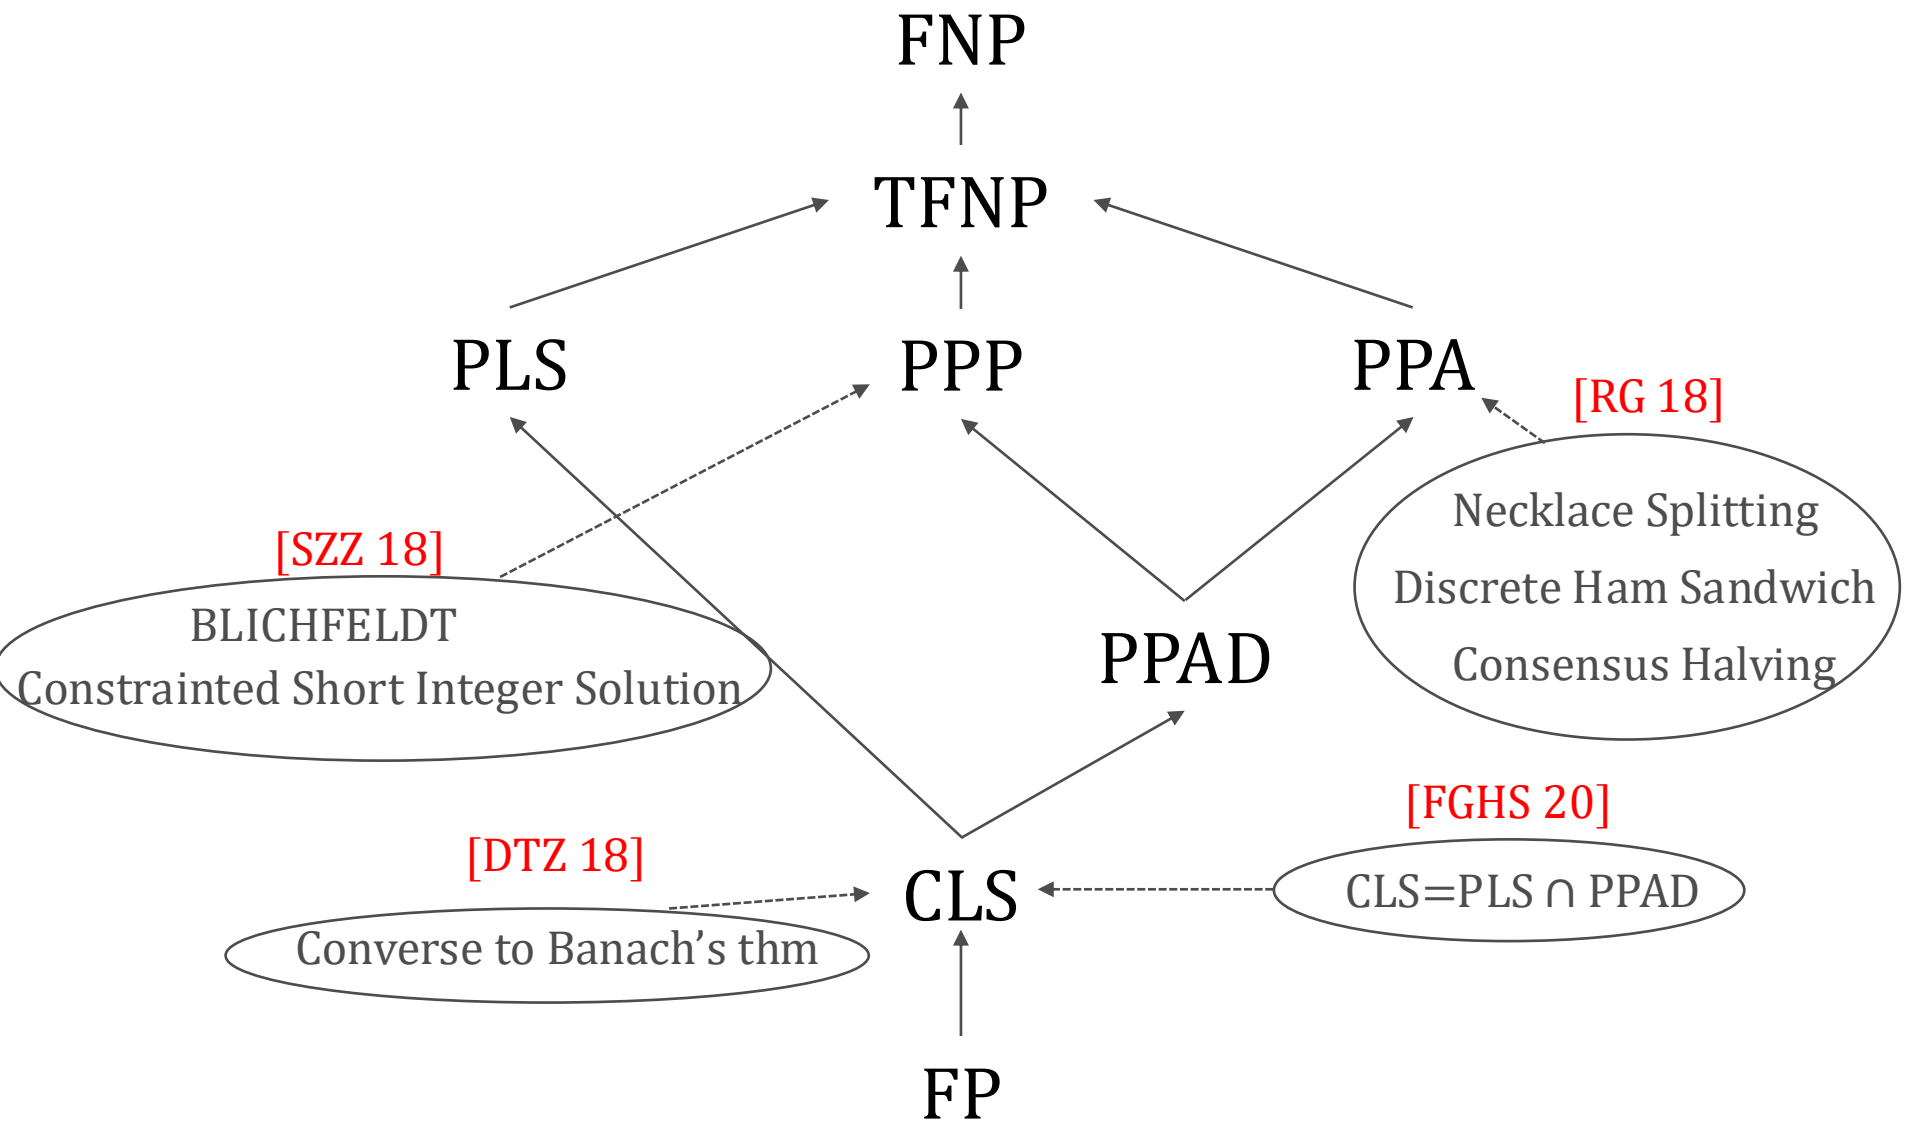
WHICH COMPLEXITY CLASS



WHICH COMPLEXITY CLASS

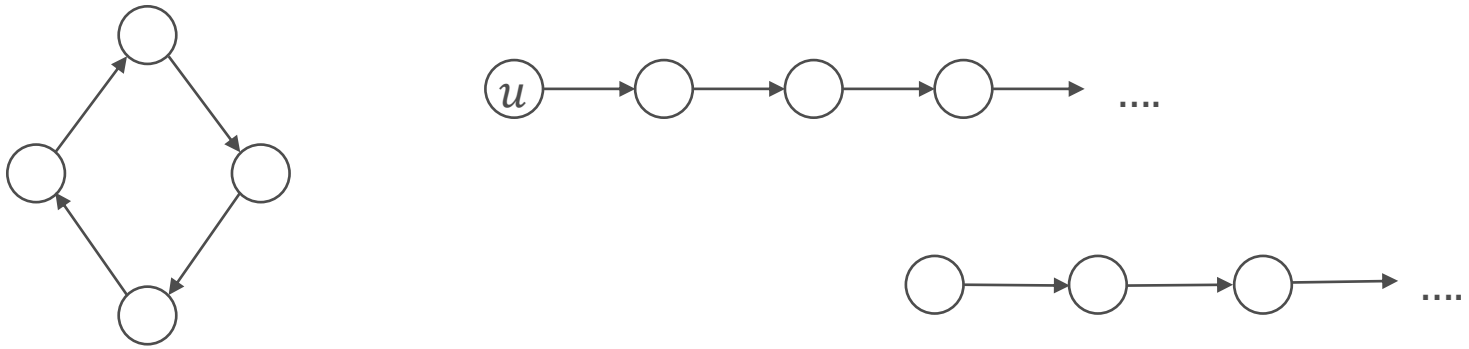


INCIDENTALLY



PPAD

- PPAD: Polynomial Parity Arguments on Directed graphs [Papadimitriou 1994]
- Input: A graph where each vertex has at most in- and out- degree at most 1. A source u.
- Goal: Find a sink or a different source!



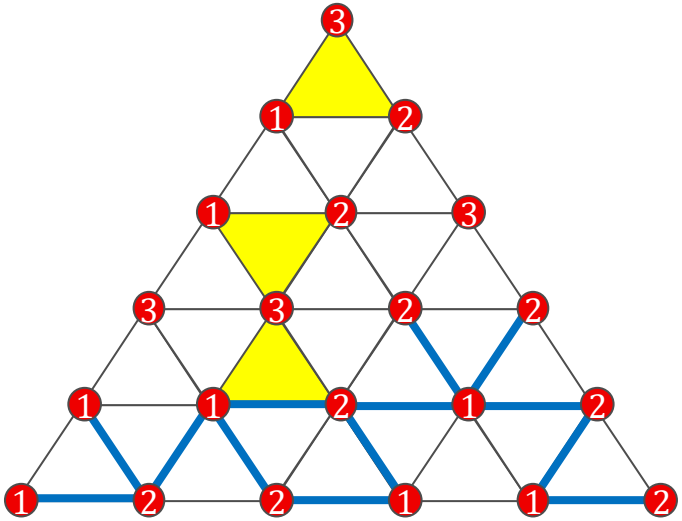
PPAD

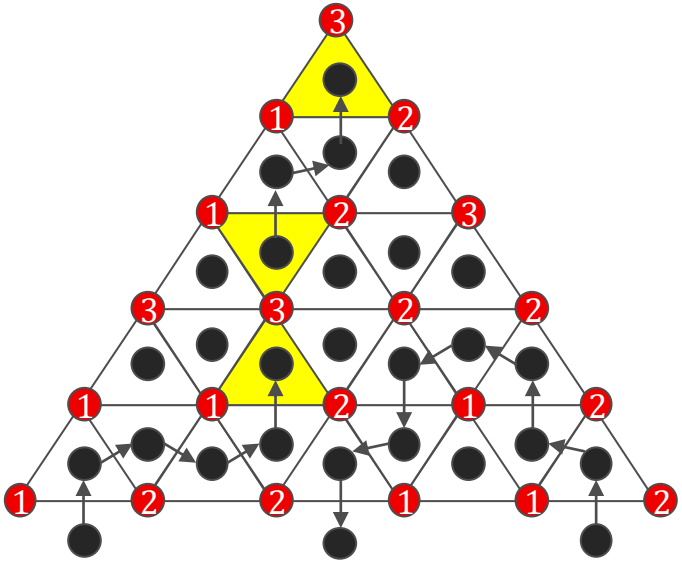
- Why not search the whole graph?
- Graph size is exponential!
- **EndOfALine**: Given two circuits S and P, with m input bits and m output bits each, such that $P(0^m) = 0^m \neq S(0^m)$, find an input $x \in \{0,1\}^m$ such that $P\big(S(x)\big) \neq$ x or $S\big(P($ 0^m.
- PPAD the set of problems reducible to EndOfALine.

WHAT DOES MNE HAVE TO DO WITH ALL THIS?

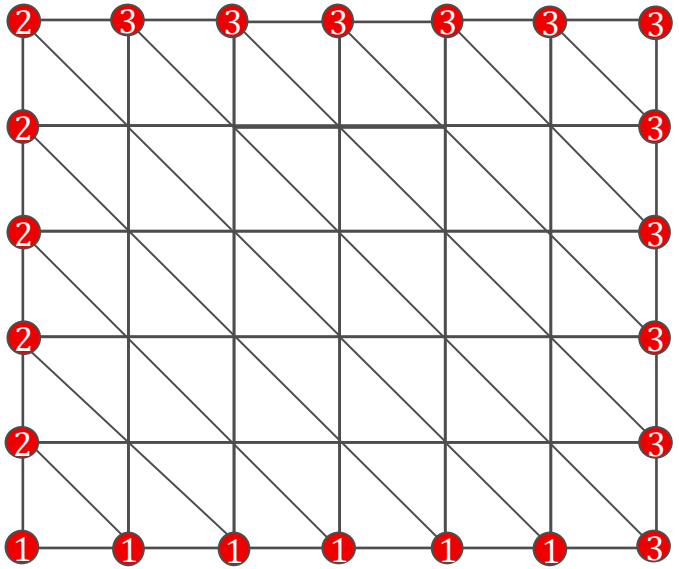
- Nash's proof that every finite game has a MNE uses a fixed point theorem argument, **Brouwer's** fixed point theorem.
- The proof of Brouwer's fixed point theorem uses **Sperner's** Lemma.
- The proof of Sperner's Lemma is at its heart an exponential time **path-following** algorithm!

SPERNER'S LEMMA



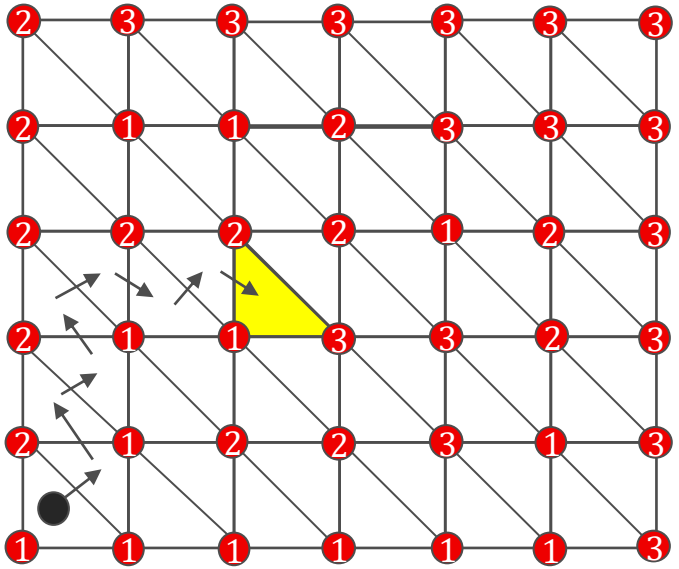


SPERNER'S LEMMA



- 2D Sperner:
  - Input: The description of a poly-time Turing machine f that gives a valid coloring. $f(p) \in \{ 0, 1, 2 \}$, where p is a node.
  - Output: A trichromatic triangle
- 2D-Sperner $\in$ PPAD
  - Obvious reduction.
- 2D-Sperner is PPAD-complete [CD 2006]

SPERNER'S LEMMA

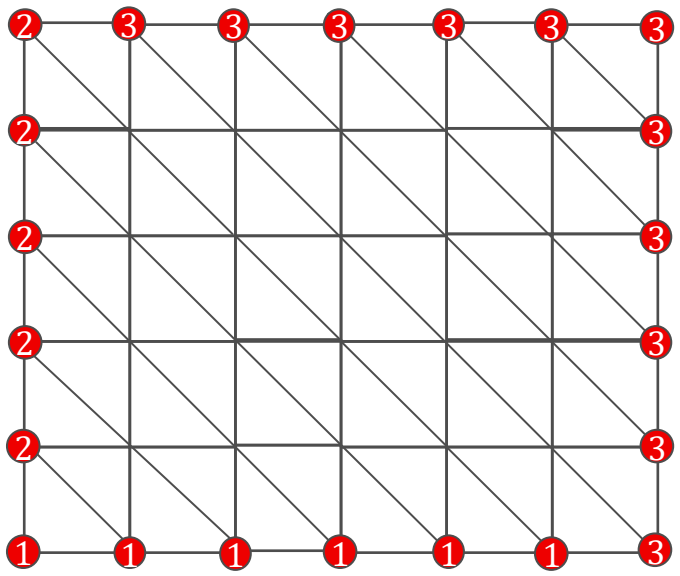


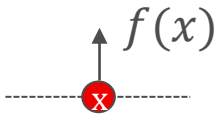
- 2D Sperner:
  - Input: The description of a poly-time Turing machine f that gives a valid coloring. $f(p) \in \{ 0, 1, 2 \}$, where p is a node.
  - Output: A trichromatic triangle
- 2D-Sperner $\in$ PPAD
  - Obvious reduction.
- 2D-Sperner is PPAD-complete [CD 2006]

BROUWER'S FIXED POINT THEOREM

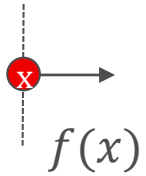
- Thm: Every continuous function f from a closed, convex and compact set C to itself has a fixed point, i.e. a point x_0 such that $f(x_0) = x_0$
- Proof (for $C = [0,1]^2$)
  - Subdivide C into tiny triangles
  - Color the edges like before.
  - For the internal nodes $x = (x_1, x_2)$:
    - If $\big(f(x)\big)_2 = f_2(x) \geq x_2$, color x with color 1
    - If $f_1(x) \geq x_1$, color x with color 2
    - If $f_1(x) \leq x_1$ and $f_2(x) \leq x_2$, color x with color 3
    - If more than 1 condition is met, pick an arbitrary color

BROUWER'S FIXED POINT THEOREM

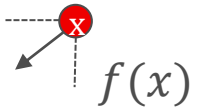




Color 1



Color 2



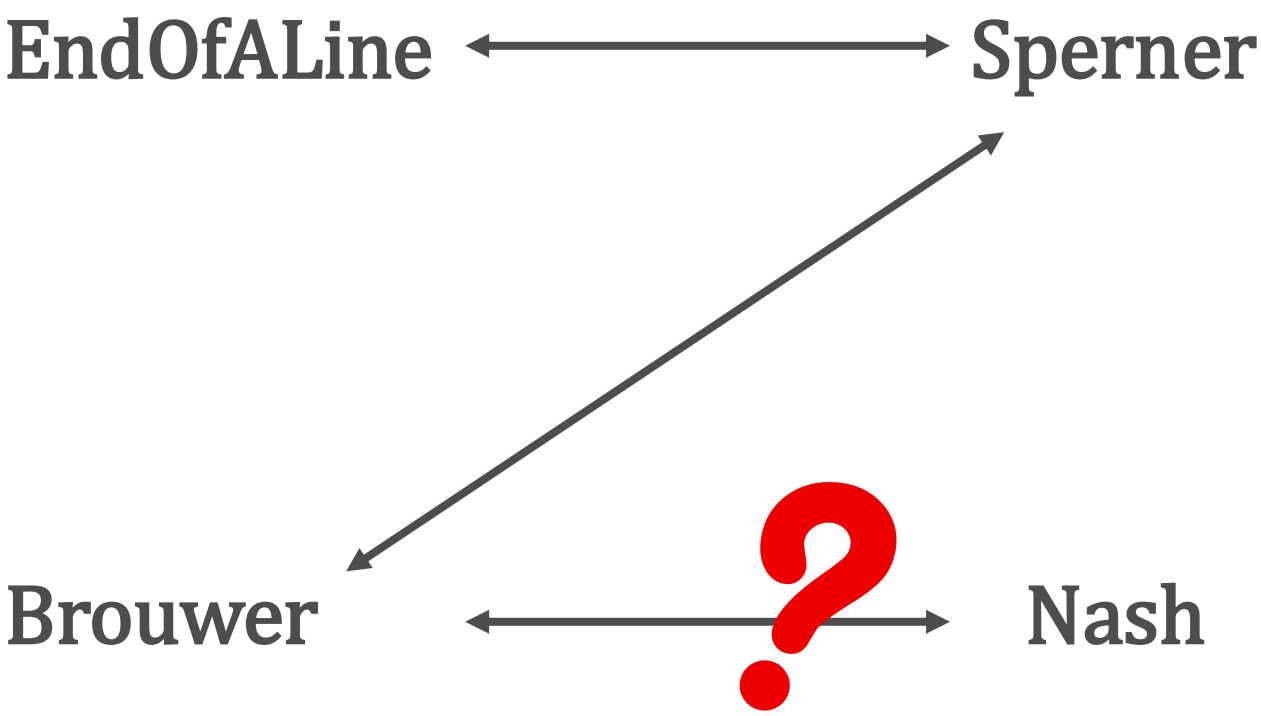
Color 3

- Color 1 = $f(x)$ farther from bottom than x
- Color 2 = $f(x)$ farther from left side than x
- Color 3 = $f(x)$ farther from top and right side than x
- Trichromatic triangle (in the limit) = $f(x)$ farther from all sides than $x = x$ is a fixed point!

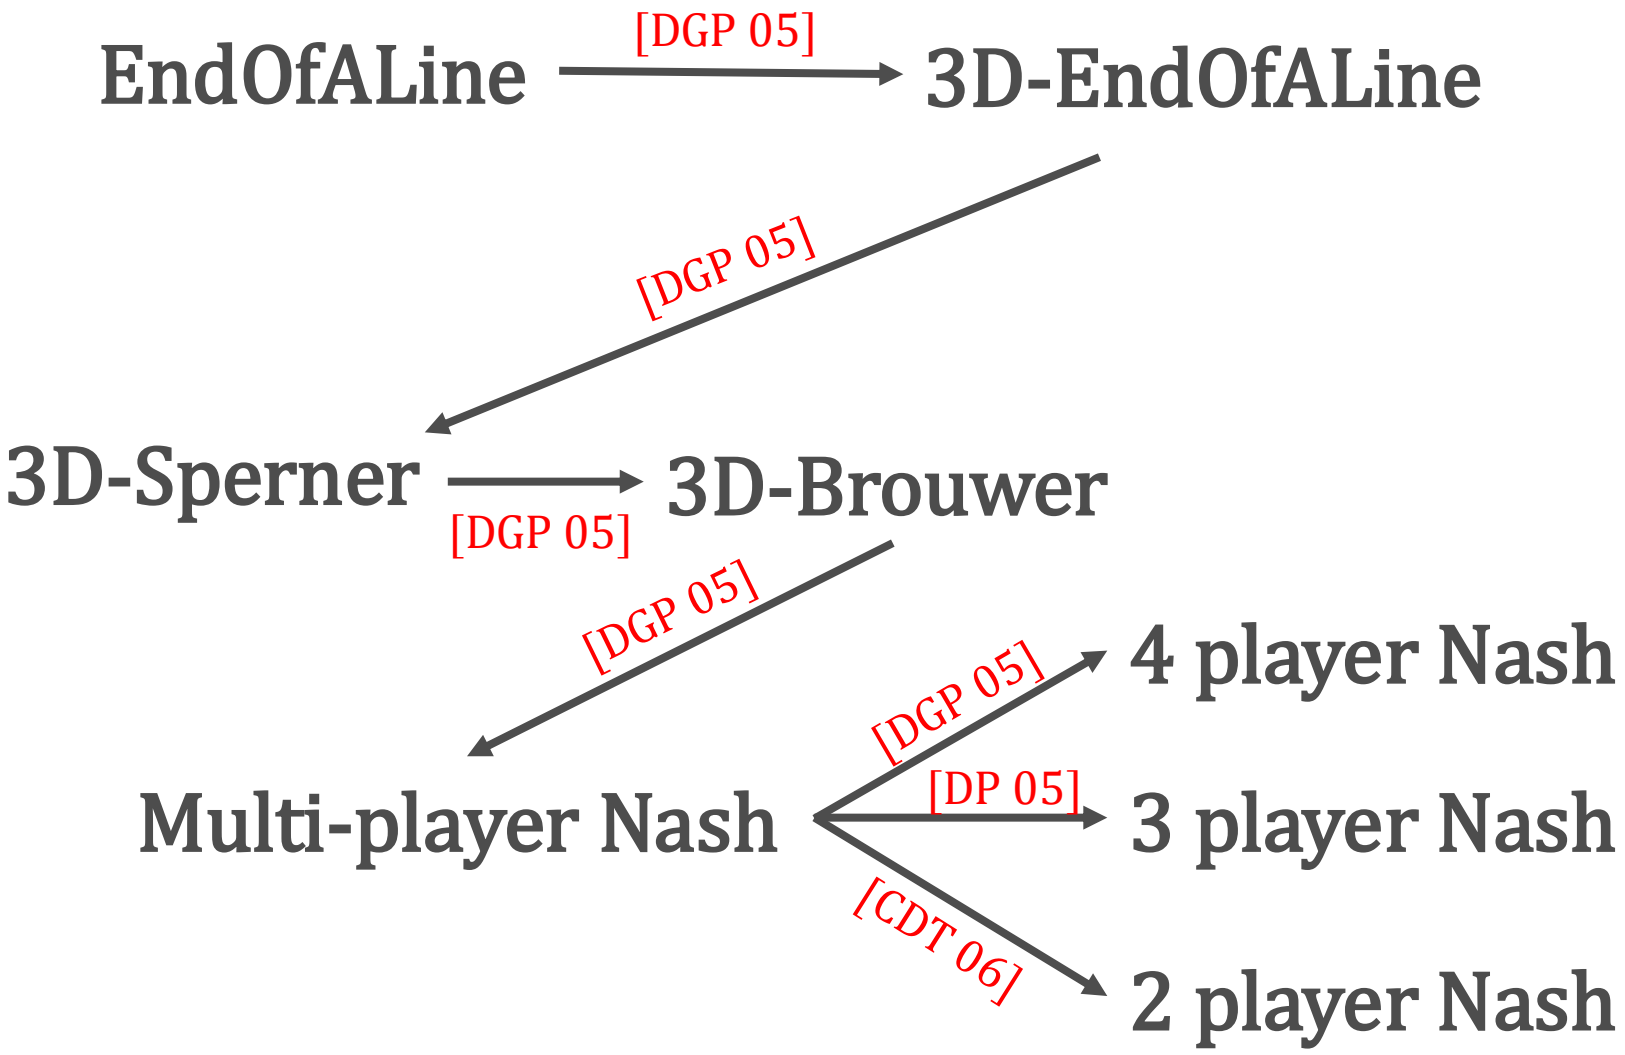
BROUWER'S FIXED POINT THEOREM

- The fixed point could be irrational!
  - We need approximation!
- Brouwer computational problem
  - Input: An algorithm A that evaluates a continuous function f from $[0,1]^n$ to $[0,1]^n$. An approximation ϵ. A Lipschitz constant c that f is claimed to satisfy.
  - Output: x such that $|f(x) - x| < \epsilon$, or a violation of the assumptions
    - $A(x)$ outside $[0,1]^n$, or $|f(x) - f(y)| > c|x - y|$
- Brouwer is PPAD-complete [DGP 05]

STORY SO FAR



THE ACTUAL STORY



THE WHOLE STORY

- Exponential approximation is PPAD complete for 3 players [DGP 06]
- Polynomial approximation is PPAD complete for 2 player NASH [CDT 06]
- Constant approximation is PPAD complete for n players [Rubinstein 15]
- Quasi-polynomial time algorithm for ϵ approximation for 2 player [LMM 03]
- Assuming ETH for PPAD, ϵ approximation takes time $2^{\Omega(n)}$ [Rubinstein 16]

REFERENCES

- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. 2009. The complexity of computing a Nash equilibrium. Commun. ACM
- Chen, X., Deng, X., and Teng, S.-H. 2009. Settling the complexity of computing two-player Nash equilibria. J. ACM
- Rubinstein, A. Inapproximability of Nash equilibrium. STOC 2015
- Rubinstein, A. Settling the Complexity of Computing Approximate Two-Player Nash Equilibria. FOCS 2016
- Lipton, R. J., Markakis, E., and Mehta, A. Playing large games using simple strategies. EC 2003



Algorithmic Game Theory

CS 58100 · Game Theory V

Alex Psomas

Fall 2026

DO PLAYERS LEARN EQUILIBRIA?

- Do we expect strategic players do converge to an equilibrium?
- Which learning algorithms (quickly) converge to an equilibrium?

BEST RESPONSE DYNAMICS

- Arguably, the most natural model of how an agent should respond is best response dynamics
- While the current outcome $\boldsymbol{s}$ is not a Pure Nash equilibrium:
  - Pick an arbitrary player i, and an arbitrary best response $s_i' \in S_i$ to the other players' current strategy s_{-i}
  - Update $\mathbf{s} \leftarrow (s_i', s_{-i})$

BEST RESPONSE DYNAMICS

- By definition, these can only converge if the game has a PNE
- But, best response dynamics might cycle in games that have PNE
- Question: When do best response dynamics converge to a PNE?

POTENTIAL GAMES

- A game is a potential game if it admits a real-valued function $\Phi: \prod_i S_i \rightarrow \mathbb{R}$ such that for all $i \in N$, for all $\boldsymbol{s} \in \prod_i S_i$, and for all $s_i' \in S_i$
  $cost_i(s_i', s_{-i}) - cost_i(\boldsymbol{s}) = \Phi(s_i', s_{-i}) - \Phi(\boldsymbol{s})$
- $Cost_i(.)$ is the cost (negative payoff) function for player i
  - Most potential games are minimization games, e.g. routing games, cost-sharing games, etc

POTENTIAL GAMES

- Theorem: In a finite potential game, from an arbitrary initial outcome, best-response dynamics converges to a PNE
- Proof:
  - In every iteration, the deviator picks a better strategy in terms of cost
  - Therefore, the value of the potential function also decreases
  - Eventually, the potential function reaches a minimum and the dynamics halts
  - This will be a PNE, by definition

NO-REGRET DYNAMICS

- More sophisticated dynamics
- Today: the model and results purely from a learning theory perspective
- Next time: what can we prove for these dynamics (and for which games)

EXAMPLE

- Suppose you have an infinite series of bits $b_1^*, b_2^*, \ldots$ being revealed one at a time
- Goal: predict the bits one by one, right before they are revealed
- There are n experts $\{e_1, \ldots, e_n\}$ you can consult
- In each round, before you guess, each expert i presents a prediction b_i^t
- Then, you decide what you want to guess

BENCHMARK

- Standard benchmark: optimal in hindsight
- Suppose that the b_i^* values can be adversarially chosen
  - That is, you predict b_i, and the adversary can now pick b_i^* to be anything
  - You won't get a single prediction correct!
- A bit too harsh!
- Instead of relaxing the adversary, we will relax the benchmark
- Goal: Compete with the best expert in hindsight

EXISTENCE OF PERFECT EXPERT

- Assume one of the experts is perfect
  - They'll never make a mistake
- Majority algorithm:
  - At time t, consult all experts who have never made a mistake in the past
  - Take the majority vote

EXISTENCE OF PERFECT EXPERT

Poll

How many mistakes does this algorithm make in T steps?

1. $\Theta(T)$
2. $\Theta(\log(T))$
3. $\Theta(n)$
4. $\Theta(logn)$



EXISTENCE OF PERFECT EXPERT

- Theorem: If there is a perfect expert, the majority algorithm makes at most $\lfloor \log_2 n \rfloor$ mistakes
- Proof:
  - $S_t = \{\ i \mid$ expert i has not made a mistake before time $t\ \}$
  - $W_t = |S_t|$
  - $W_0 = n$
  - If the majority makes a mistake at time t:
    - total number of mistakes increases by 1
    - $W_{t+1} \leq \lfloor \frac{W_t}{2} \rfloor$
  - $W_t \geq 1$ for all t, since a perfect expert exists
  - Total number of mistakes at most $\lfloor \log_2 n \rfloor$

NO REGRET

- Goal: relax the assumption that a perfect expert exists
- We'll be able to get positive results in a more general model

MODEL

- T time steps
- Set A of $n \geq 2$ actions (the experts)
- At time $t = 1, \ldots, T$:
  - Pick a mixed strategy/distribution over actions p^t
  - The adversary picks a cost vector $c^t : A \to [0,1]$
    - The only important thing is the costs are bounded
  - An action a^t is drawn according to p^t and we pay cost $c^t(a^t)$
- Goal: Compete with the best action in hindsight

REGRET

- Definition: The regret of the action sequence $a^1, \ldots, a^T$ with respect to action a is

$$\frac{1}{T}\left(\sum_{t=1}^{T} c^t(a^t) - \sum_{t=1}^{T} c^t(a)\right)$$

- An algorithm has no regret if the expected regret with respect to every action $a \in A$ is $o(1)$, as T goes to infinity
  - The expectation is with respect to the randomness in the algorithm and any randomness in the adversary
  - Perhaps better name is vanishing regret

RANDOMIZATION IS NECESSARY

- Claim: If a no-regret algorithm exists, it must be randomized
  - Consider any deterministic algorithm
  - At each time step t the algorithm commits to $a^t \in A$
  - Adversary sets $c^t(a^t) = 1$ and $c^t(a) = 0$ for all $a \neq a^t$
  - After T steps, the algorithm has paid T
  - There is some action with cost at most T/n
  - The regret is $\approx 1 - 1/n$, which does not go to zero as T increases

NO-REGRET

- Theorem: There exist (simple) no-regret algorithms, with expected regret $O(\sqrt{lnn/T})$ with respect to every fixed action

THE MULTIPLICATE WEIGHTS ALGORITHM

- Initialize $w^1(a) = 1$ for all $a \in A$
- For $t = 1, \ldots, T$:
  - Play action a with probability $p^t(a) = \frac{w^t(a)}{\sum_{a' \in A} w^t(a')}$
  - After cost vector c^t is revealed update weights according to

$$w^{t+1}(a) = w^t(a) \cdot (1 - \epsilon)^{c^t(a)}$$

THE MULTIPLICATE WEIGHTS ALGORITHM

- $w^{t+1}(a) = w^t(a) \cdot (1-\epsilon)^{c^t(a)}$
- Intuition:
  - $c^t(a) \in \{0,1\}$: weight either stays the same, or gets multiplied by $1-\epsilon$
  - $\epsilon \to 0$: closer to uniform distribution
  - $\epsilon \to 1$: closer to "pick best so far"

THE MULTIPLICATE WEIGHTS ALGORITHM: ANALYSIS

- $OPT = \sum_{t=1}^{T} c^t(a^*)$
- Plan:
  - Step 1: relate total weight with OPT
  - Step 2: relate total weight with cost of MW
- Let $\Gamma^t = \sum_{a \in A} w^t(a)$ be the total weight in step t
- Step 1:
  - $\Gamma^T \geq w^T(a^*)$
  - $= w^1(a^*) \prod_{t=1}^{T} (1-\epsilon)^{c^t(a^*)}$
  - $= (1-\epsilon)^{OPT}$

THE MULTIPLICATE WEIGHTS ALGORITHM: ANALYSIS

- Step 2:
  - Expected cost v^t of MW at time t

$$v^t = \sum_{a \in A} p^t(a) \cdot c^t(a) = \sum_{a \in A} \frac{w^t(a)}{\Gamma^t} \cdot c^t(a)$$

  - $\Gamma^{t+1} = \sum_{a \in A} w^{t+1}(a)$
  - $= \sum_{a \in A} w^t(a) \cdot (1-\epsilon)^{c^t(a)}$
  - $\leq \sum_{a \in A} w^t(a) \cdot \left(1 - \epsilon c^t(a)\right)$
  - $= \Gamma^t \cdot (1 - \epsilon v^t)$

$(1-\epsilon)^x \leq 1 - \epsilon x,$
$\forall \epsilon \in [0, ½], x \in [0,1]$

THE MULTIPLICATE WEIGHTS ALGORITHM: ANALYSIS

- Goal: bound $\sum_{t=1}^{T} v^t$
- What we've shown
  - $(1-\epsilon)^{OPT} \leq \Gamma^T \leq \Gamma^1 \cdot \prod_{t=1}^{T}(1-\epsilon v^t)$
- Therefore

$$OPT \cdot \ln(1-\epsilon) \leq \ln(n) + \sum_{t=1}^{T} \ln(1-\epsilon v^t)$$

- Taylor expansion:

$$\ln(1-x) = -x - \frac{x^2}{2} - \frac{x^3}{3} - \cdots$$

  - $\ln(1-x) \leq -x$
- Also, $1 - x \geq e^{-x\,-x^2}$, for $x \geq \frac{1}{2}$
  - $\ln(1-x) \geq -x - x^2$

THE MULTIPLICATE WEIGHTS ALGORITHM: ANALYSIS

- Goal: bound $\sum_{t=1}^{T} v^t$
- Putting it back

$$-OPT \cdot (\epsilon + \epsilon^2) \leq \ln(n) + \sum_{t=1}^{T} -\epsilon v^t$$

- Re-arranging:

$$\sum_{t=1}^{T} v^t \leq OPT \cdot (1 + \epsilon) + \frac{\ln(n)}{\epsilon}$$

- Or

$$\sum_{t=1}^{T} v^t \leq OPT + \epsilon T + \frac{\ln(n)}{\epsilon}$$

THE MULTIPLICATE WEIGHTS ALGORITHM: ANALYSIS

$$\sum_{t=1}^{T} v^t \leq OPT + \epsilon T + \frac{\ln(n)}{\epsilon}$$

- Set ϵ to equalize the two terms

$$\epsilon = \sqrt{\ln(n)\ /T}$$

- Overall:

$$\sum_{t=1}^{T} v^t \leq OPT + 2\sqrt{Tln(n)}$$

THE MULTIPLICATE WEIGHTS ALGORITHM: ANALYSIS

- That was cheating!
  - ϵ was dependent of T!
- Easy fix:
  - At day t pretend T is the smallest power of 2 bigger than t
  - Analysis carries over

RECAP

- Best response dynamics
  - Converge in potential games
- No-regret algorithms
  - Multiplicative weight update algorithm



Algorithmic Game Theory

CS 58100 · Game Theory VI

Alex Psomas

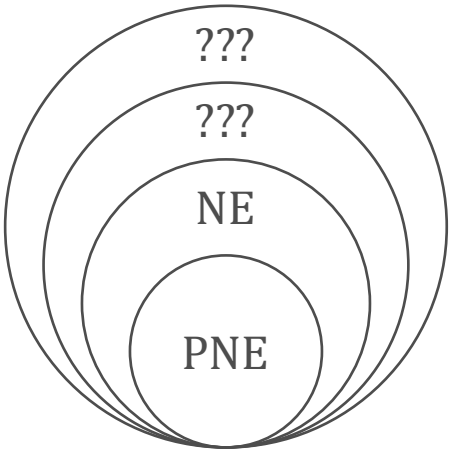
Fall 2026

STORY SO FAR

- Nash equilibria
  - Always exist 😀
  - Hard to compute 😩
- Pure Nash equilibria
  - Very stable as an equilibrium notion 😀
  - Might not exist 😩

TODAY

- More equilibrium concepts



CORRELATED EQUILIBRIUM

- For now $N = \{1,2\}$
- A mediator chooses a pair of strategies (s_1, s_2) according to a know distribution p over $S_1 \times S_2$
- Mediator reveals s_1 to player 1 and s_2 to player 2
  - Player 1's perspective having seen s_1

$$\Pr[s_2|s_1] = \frac{\Pr[s_1 \,\&\, s_2]}{\Pr[s_1]} = \frac{p(s_1, s_2)}{\Pr[s_1]}$$

CORRELATED EQUILIBRIUM

- Player 1 is best responding if, for all $s_1' \in S_1$:

$$\sum_{s_2 \in S_2} \Pr[s_2|s_1]\, u_1(s_1, s_2) \geq \sum_{s_2 \in S_2} \Pr[s_2|s_1]\, u_1(s_1', s_2)$$

- Player 2 is best responding if, for all $s_2' \in S_2$:

$$\sum_{s_1 \in S_1} \Pr[s_1|s_2]\, u_2(s_1, s_2) \geq \sum_{s_1 \in S_1} \Pr[s_1|s_2]\, u_2(s_1, s_2')$$

- Definition: p is a correlated equilibrium (CE) if all agents are best responding

EXAMPLE
GAME OF TRAFFIC

	Stop	Go
Stop	3,3	1,4
Go	4,1	0,0

- Should you stop at an intersection with no STOP signs?
  - Aka game of chicken
- Worst case scenario: both go
- Ideal scenario: you stop, other person goes

EXAMPLE
GAME OF TRAFFIC

	Stop	Go
Stop	3,3	1,4
Go	4,1	0,0

- PNEs: (Stop, Go) and (Go, Stop)
  - Social welfare (sum of utilities) is 5
- NE: $\left(\left(\frac{1}{2},\frac{1}{2}\right),\left(\frac{1}{2},\frac{1}{2}\right)\right)$
  - Social welfare is 4
- Can we do better?

EXAMPLE
GAME OF TRAFFIC

- Mediator (traffic light):
  - With probability ½ recommend (Stop, Go)
  - With probability ½ recommend (Go, Stop)
- Player 1's perspective:
  - Conditioned on being recommended Stop, player 2 got recommended Go.
  - "Assuming that player 2 follows this recommendation, I better Stop"
- Social welfare: 5
- Can we do better?

	Stop	Go
Stop	3,3	1,4
Go	4,1	0,0

EXAMPLE
GAME OF TRAFFIC

	Stop	Go
Stop	3,3	1,4
Go	4,1	0,0

- Mediator (fancy traffic light):
  - (Stop, Go): 1/3
  - (Go, Stop): 1/3
  - (Stop, Stop): 1/3
- This is also a CE
- Social welfare:
  - $6 \cdot \frac{1}{3} + 5 \cdot \frac{1}{3} + 5 \cdot \frac{1}{3} = \frac{16}{3}$

FANCY MEDIATOR

- Instead of a mediator, you can also use a hat
- Balls in the hat are labelled with “Stop” and “Go”
- Each player takes a ball (blindfolded)

Poll 1

Which composition of balls implements the previous distribution?



1. 1 Stop, 1 Go
2. 1 Stop, 2 Go
3. 2 Stop, 1 Go
4. 2 Stop, 2 Go

FORMAL DEFINITION

- Distributions $\vec{\sigma} = \sigma_1, \ldots, \sigma_n$ over strategy sets $S_1, \ldots, S_n$ constitute a Nash equilibrium (NE) if, for every player i, and every unilateral deviation $s_i' \in S_i$

$$\mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i, s_{-i})] \geq \mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i', s_{-i})]$$

- A distribution $\vec{\sigma}$ over strategy profiles $\mathcal{S} = \prod_i S_i$ constitutes a correlated equilibrium (CE) if, for every player i, strategy s_i, and every deviation $s_i' \in S_i$

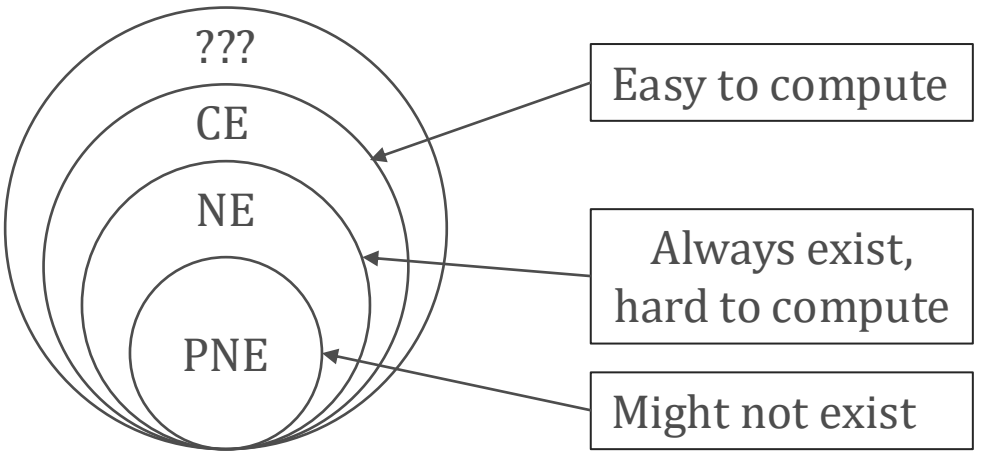
$$\mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i, s_{-i})|s_i] \geq \mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i', s_{-i})|s_i]$$

LINEAR PROGRAM

- Correlated equilibria can be computed via linear programming
- Variables: $p(s_1, s_2)$
- Constraints:
  - $\sum_{s_2 \in S_2} p(s_1, s_2)\, u_1(s_1, s_2) \geq \sum_{s_2 \in S_2} p(s_1, s_2)\, u_1(s_1', s_2), \forall s_1, s_1'$
  - $\sum_{s_1 \in S_1} p(s_1, s_2)\, u_2(s_1, s_2) \geq \sum_{s_1 \in S_1} p(s_1, s_2)\, u_2(s_1, s_2'), \forall s_2, s_2'$
  - $\sum_{s_1 \in S_1} \sum_{s_2 \in S_2} p(s_1, s_2) = 1$
  - $p(s_1, s_2) \geq 0, \forall s_1, s_2$

WHAT'S COMING

- Every Nash equilibrium is a Correlated equilibrium, but not the other way
- Unlike Nash equilibria, Correlated Equilibria can be computed in polynomial time
  - There are a couple of ways to do this



WHAT MORE DO YOU WANT?

- Correlated equilibria always exist (since Nash equilibria always exist) and can be computed efficiently!
- Let's be a little greedy!
- We will expand our notion of equilibrium, and manage to get something that we can compute even with simple, distributed learning algorithms

COARSE CORRELATED EQUILIBRIUM

- A distribution $\vec{\sigma}$ over strategy profiles $\mathcal{S} = \prod_i S_i$ constitutes a correlated equilibrium (CE) if, for every player i, strategy s_i, and every deviation $s_i' \in S_i$

$$\mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i, s_{-i})|s_i] \geq \mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i', s_{-i})|s_i]$$

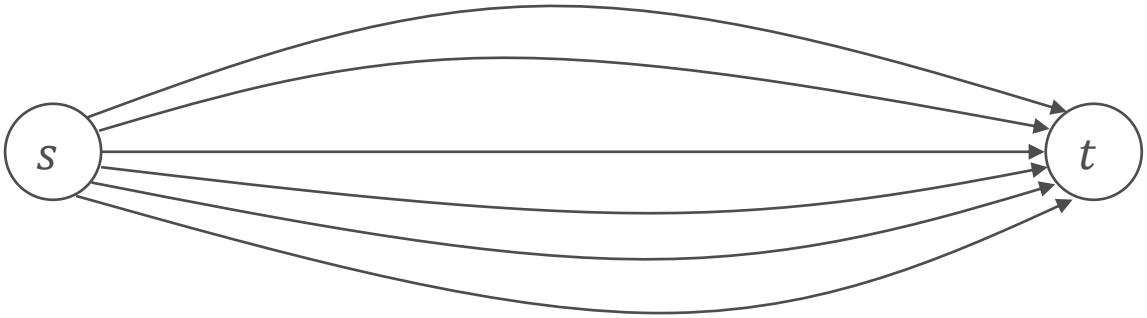
- A distribution $\vec{\sigma}$ over strategy profiles $\mathcal{S} = \prod_i S_i$ constitutes a coarse correlated equilibrium (CCE) if, for every player i, and every deviation $s_i' \in S_i$

$$\mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i, s_{-i})] \geq \mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i', s_{-i})]$$

COARSE CORRELATED EQUILIBRIUM

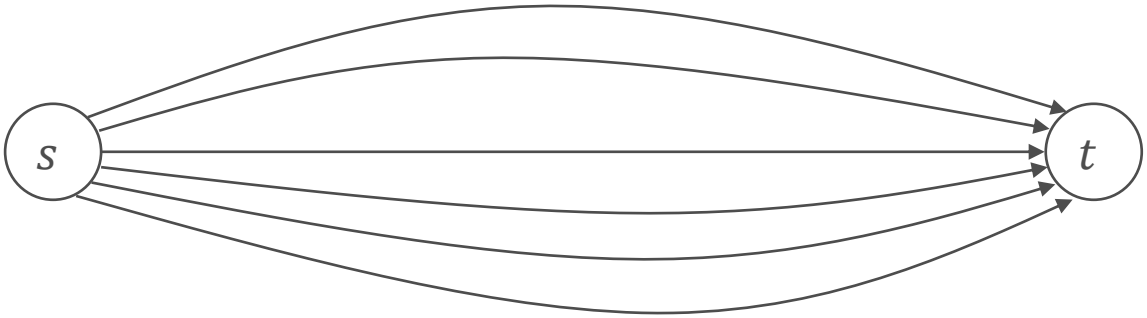
- In a CCE, when player i considers s_i' she knows only the joint distribution σ and not the realization s_i
- In a CCE we are "protected" only against unilateral, unconditional deviations
- Every CE is a CCE, but not the other way around

EXAMPLE



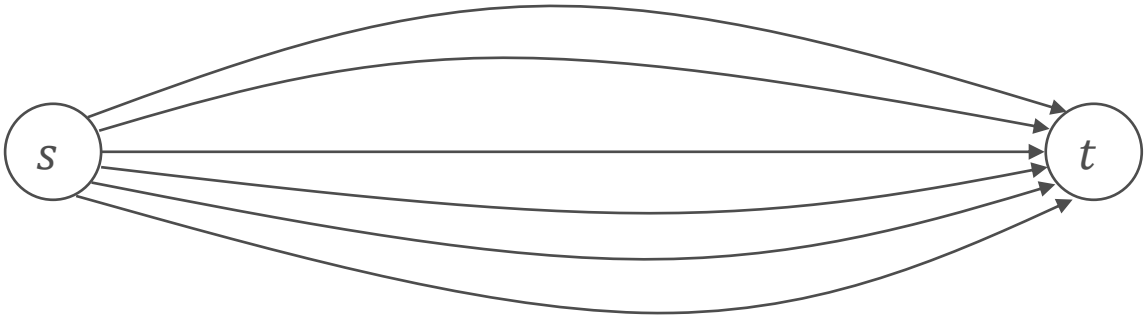
- Common source s, common sink t
- 6 parallel edges
  - Cost of each edge is $c(x) = x$
- Four players want to route one unit of flow (each) from s to t

EXAMPLE



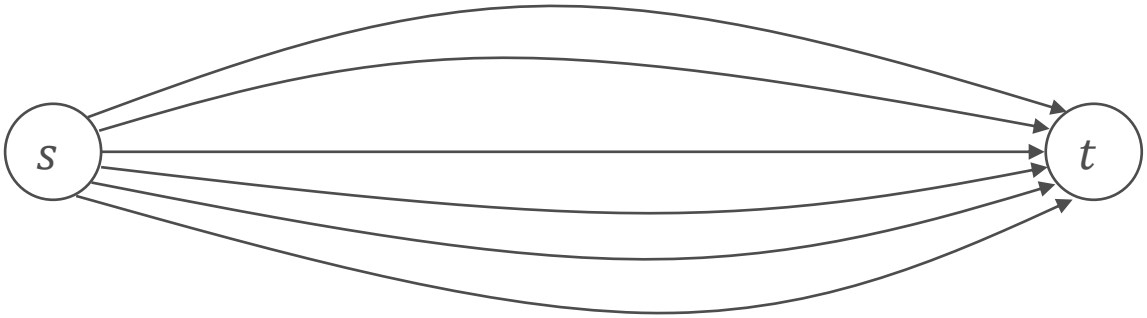
- PNEs:
  - $\binom{6}{4}$ outcomes where each player picks different edge

EXAMPLE



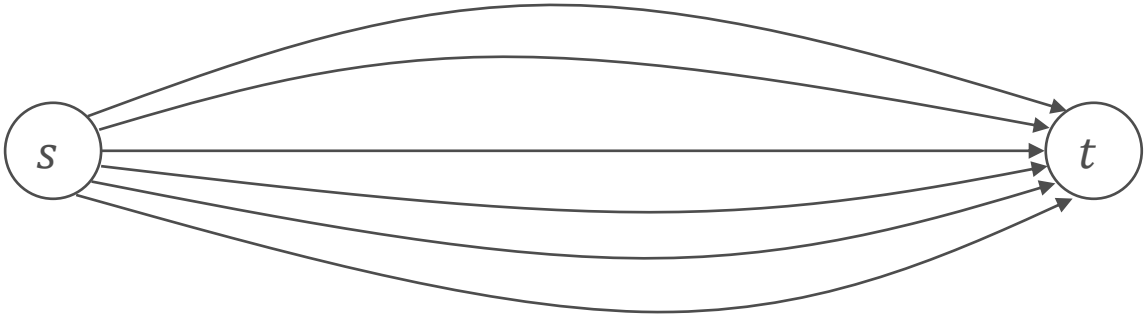
- MNE:
  - Each player chooses an edge uniformly at random

EXAMPLE



- CE:
  - Uniform distribution over all outcomes in which there is one edge with two players and two edges with one player

EXAMPLE



- CCE:
  - Uniform distribution over outcomes in which there is one edge with two players and two edges with one player, but the edges chosen are either {0,2,4} or {1,3,5}
    - This is not a CE, since conditioned on being recommended even you want to switch to odd, and vice versa

NO REGRET LEARNING AND CCE

- A distribution $\vec{\sigma}$ over strategy profiles $\mathcal{S} = \prod_i S_i$ constitutes a coarse correlated equilibrium (CCE) if, for every player i, and every deviation $s_i' \in S_i$

$$\mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i, s_{-i})] \geq \mathbb{E}_{\vec{s}\sim\vec{\sigma}}[u_i(s_i', s_{-i})]$$

- Theorem: There exist (simple) no-regret algorithms, with expected regret $O(\sqrt{lnn/T})$ with respect to every fixed action

NO REGRET LEARNING AND CCE

- A distribution $\vec{\sigma}$ over strategy profiles $\mathcal{S} = \prod_i S_i$ constitutes a coarse correlated equilibrium (CCE) if, for every player i, and every deviation $s_i' \in S_i$

$$\mathbb{E}_{\vec{s}\sim\vec{\sigma}}\left[c_i(s_i, s_{-i})\right] \leq \mathbb{E}_{\vec{s}\sim\vec{\sigma}}\left[c_i(s_i', s_{-i})\right]$$

- Theorem: There exist (simple) no-regret algorithms, with expected regret $O(\sqrt{lnn/T})$ with respect to every fixed action
  - There are obvious analogs for payoff maximization (vs cost minimization) games

NO REGRET LEARNING AND CCE

- In each time step $t = 1, \dots, T$
  - Each player i simultaneously and independently chooses a mixed strategy p_i^t using a no-regret algorithm
    - This selects (pure) strategy $s_{i,j}$ with probability $p_{i,j}^t$
  - Each player i receives a cost vector c_i^t, where $c_i^t(s_i) = E_{s_{-i} \sim p_{-i}^t}[c_i(s_i, s_{-i})]$

- No-regret guarantee for player i implies that

$$\sum_{t=1}^{T} E_{s_i \sim p_i^t}\left[c_i^t(s_i)\right] \leq \sum_{t=1}^{T} c_i^t(s_i') + 2\sqrt{T ln(n)}$$

  - Note: c_i^t was already defined as an expectation

NO REGRET LEARNING AND CCE

- Let $\sigma^t = \prod_{i=1}^{n} p_i^t$ denote the outcome distribution at time t
- Let $\sigma = \frac{1}{T}\sum_{t=1}^{T} \sigma^t$ denote the time-averaged history of these distributions
- Theorem: σ is an approximate CCE. That is

$$E_{\boldsymbol{s}\sim\sigma}[c_i(\boldsymbol{s})] \leq E_{\boldsymbol{s}\sim\sigma}[c_i(s_i', s_{-i})] + \epsilon$$

where ϵ is an upper bound on the regret of every player after T rounds

  - The σ^ts are not a CCE, but looking at the "average" game so far, we have a CCE

PROOF

- By definition, for each player i

$$E_{\boldsymbol{s}\sim\sigma}[c_i(\boldsymbol{s})] = \frac{1}{T}\sum_{t=1}^{T} E_{\boldsymbol{s}\sim\sigma^t}[c_i(\boldsymbol{s})]$$

$$E_{\boldsymbol{s}\sim\sigma}[c_i(s_i', s_{-i})] = \frac{1}{T}\sum_{t=1}^{T} E_{\boldsymbol{s}\sim\sigma^t}[c_i(s_i', s_{-i})]$$

- The first RHS is the time-average cost when playing according to a no-regret algorithm
- The second RHS is the time-average cost when playing a fixed strategy s_i'
- For any no-regret algorithm, the second RHS if the first plus ϵ

WHAT ABOUT CORELATED EQUILIBRIA?

- Can we define a learning procedure that converges to a correlated equilibrium?

SWAP-REGRET

Definition: An online decision-making algorithm has no swap regret if for every adversary, the expected swap-regret

$$\frac{1}{T}\left(\sum_{t=1}^{T} c^t(a^t) - \sum_{t=1}^{T} c^t\left(\delta(a^t)\right)\right)$$

with respect to every switching function $\delta: A \to A$, is $o(1)$ as $T \to \infty$

SWAP-REGRET

Theorem: Each player i simultaneously and independently chooses a mixed strategy p_i^t using a no-swap-regret algorithm. Let $\sigma^t = \prod_{i=1}^{n} p_i^t$ denote the outcome distribution at time t. Let $\sigma = \frac{1}{T}\sum_{t=1}^{T} \sigma^t$ denote the time-averaged history of these distributions. Then σ is an $\epsilon -$ correlated equilibrium:

$$E_{s\sim\sigma}[c_i(s)] \leq E_{s\sim\sigma}[c_i(\delta(s_i), s_{-i})] + \epsilon$$

for every player i and switching function $\delta: S_i \rightarrow S_i$

SWAP-REGRET

- **Theorem**([Blum, Mansour 2007): If there is a no-regret algorithm, there is a no-swap regret algorithm.

Proof intuition:

- Initialize n copies of MW (one for each expert)!
- We will get the recommendation of each copy, and find some way to aggregate them into one decision
- Let S_j be the set of time steps in which your algorithm picked expert j
- No regret with respect to j on S_j

$$\frac{1}{|S_j|}\sum_{t\in S_j} c^t(j) \leq \frac{1}{|S_j|}\min_i \sum_{t\in S_j} c^t(i) + o(|S_j|)$$

implies no swap regret

TAKE HOME MESSAGES

- There are simple notions of equilibria beyond Nash equilibria
- Very simple learning algorithms can be used to quickly approximate coarse correlated equilibria
- Almost the same thing is true for correlated equilibria (the dynamics are a bit more complicated)

REFERENCES

- Tim Roughgarden's lecture notes:

http://timroughgarden.org/f13/l/l17.pdf

http://timroughgarden.org/f13/l/l18.pdf



Algorithmic Game Theory

CS 58100 · Game Theory VII

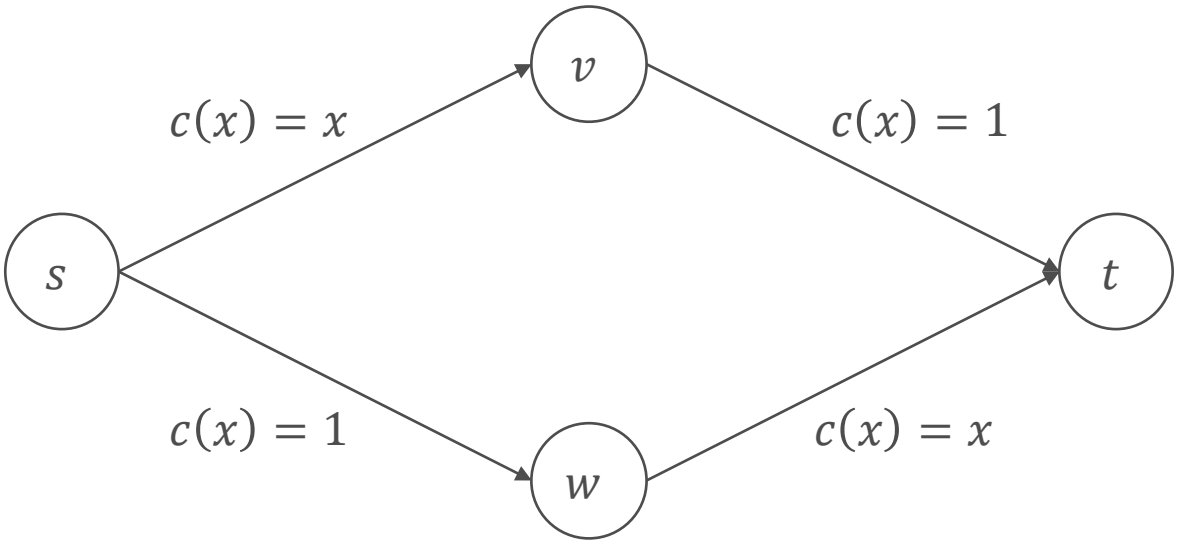
Alex Psomas

Fall 2026

FINAL EXAM

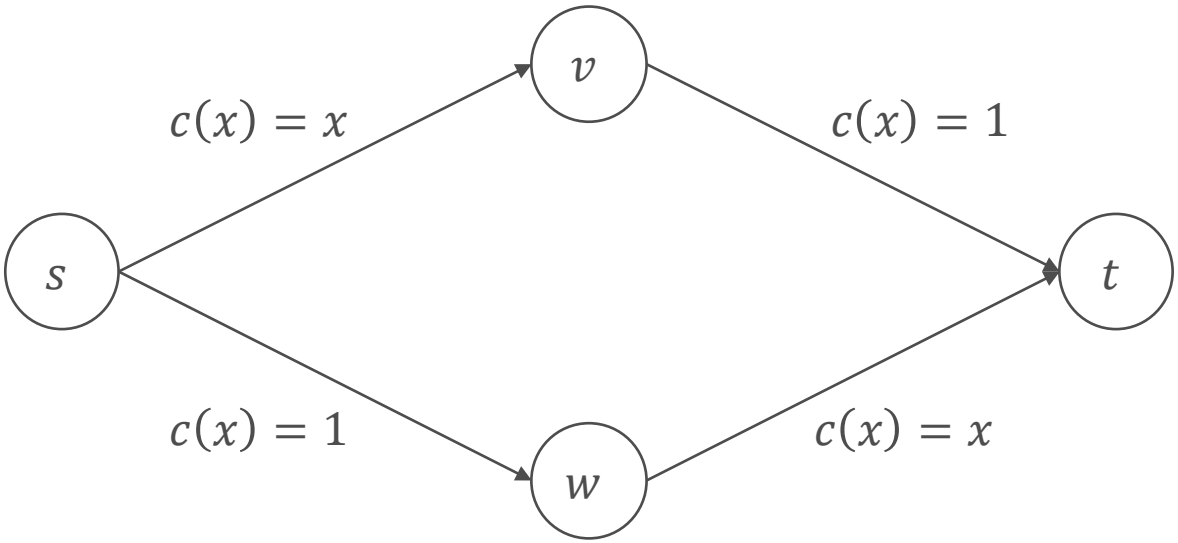
- We have a date!
- Wed, 12/16, 08:00am - 10:00am
- WALC 2007 90 45

BRAESS'S PARADOX



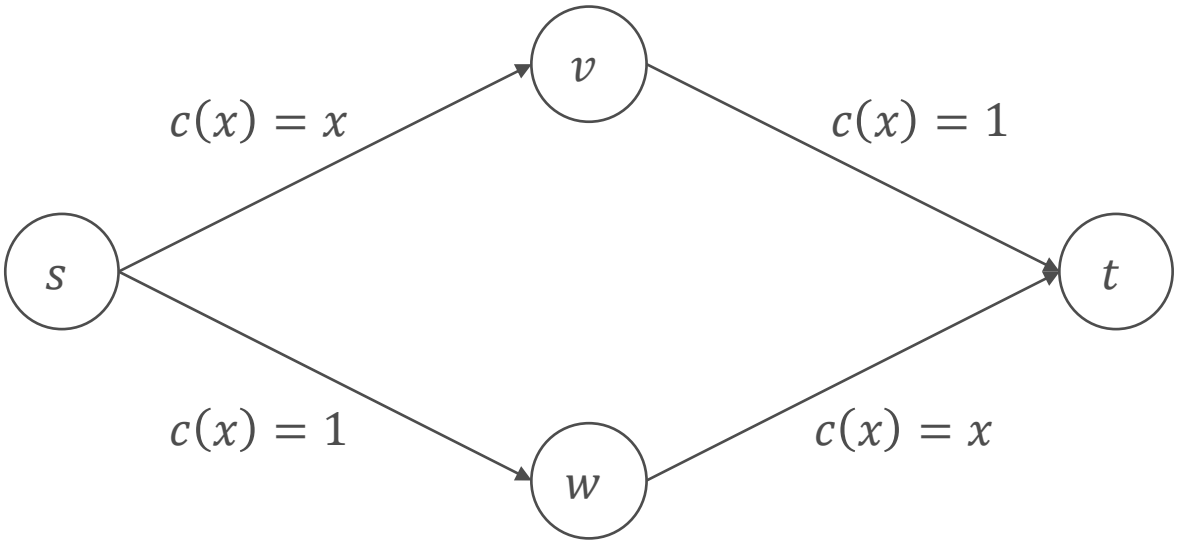
- One unit of "selfish" flow will go from s to t
- What will happen?

BRAESS'S PARADOX



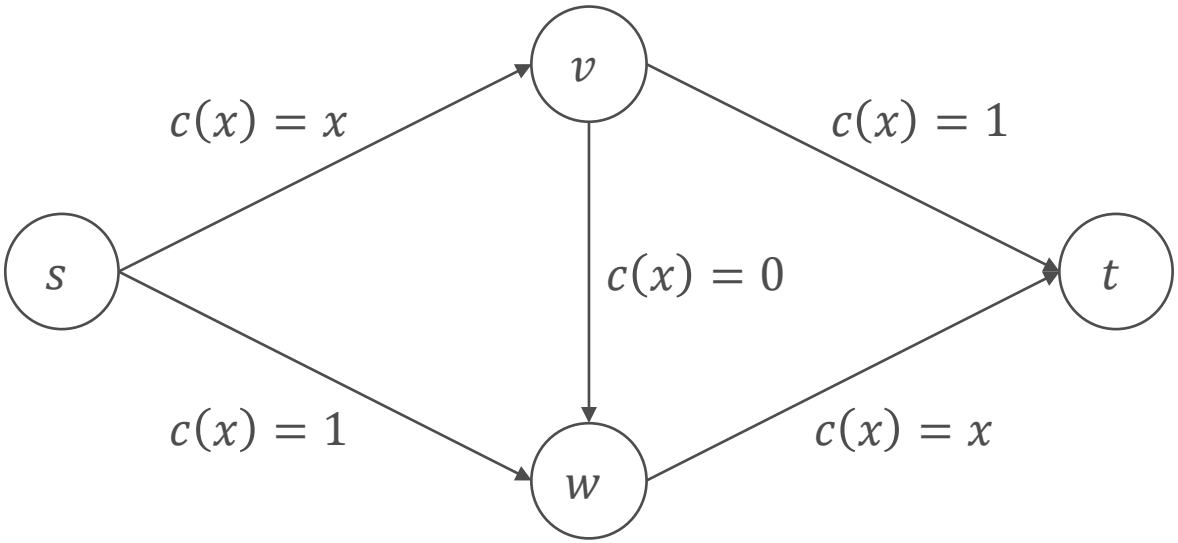
- If an x fraction uses $s \to v \to t$, then they get to t after $1 + x$ hours
- The remaining $1 - x$ fraction uses $s \to w \to t$ and gets to t in $1 + (1 - x)$ hours

BRAESS'S PARADOX



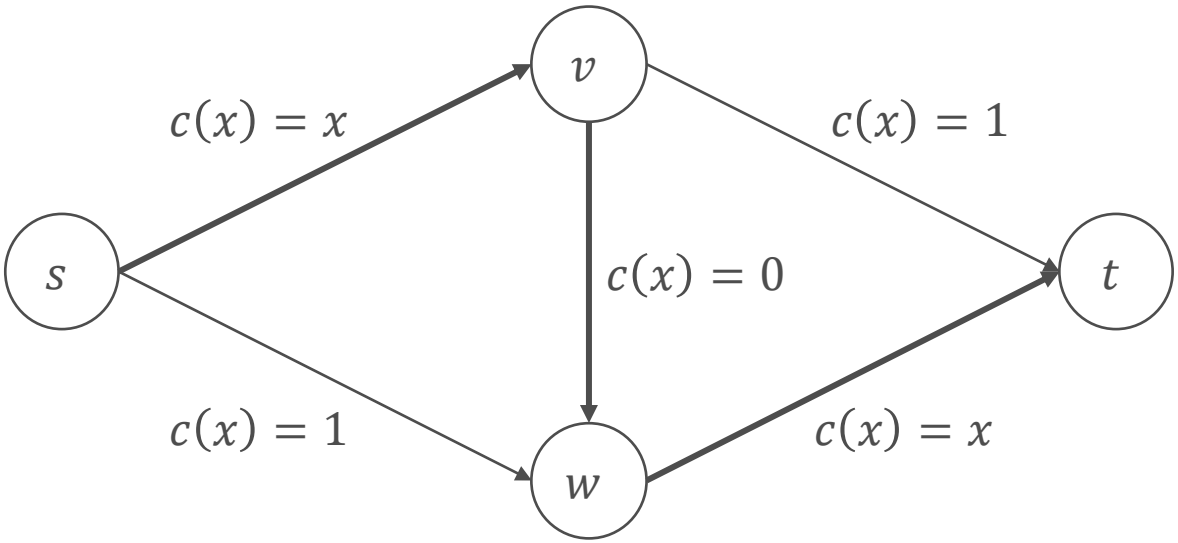
- Splitting equally is a Nash equilibrium
- Total commute time $\left(1 + \frac{1}{2}\right) = \frac{3}{2}$

BRAESS'S PARADOX



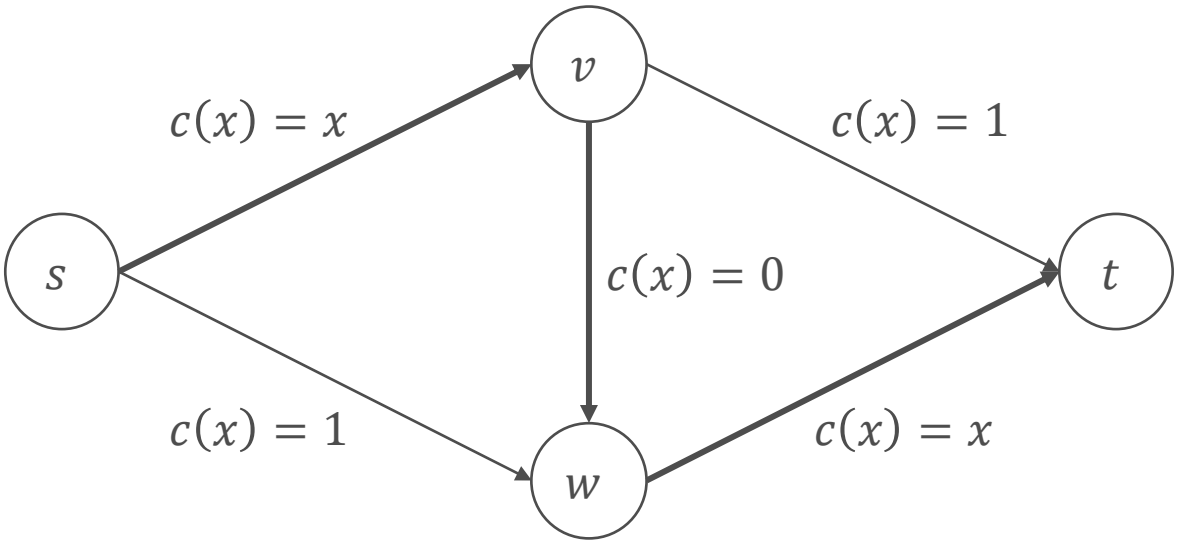
- Government invests in a new, really fast road
- What now?

BRAESS'S PARADOX



- The path $s \to v \to w \to t$ is never worse than $s \to v \to t$ or $s \to w \to t$
- Therefore, no matter what everyone else is doing, you want to take $s \to v \to w \to t$

BRAESS'S PARADOX

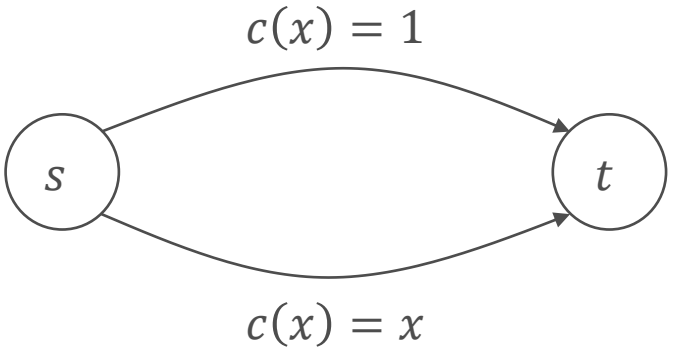


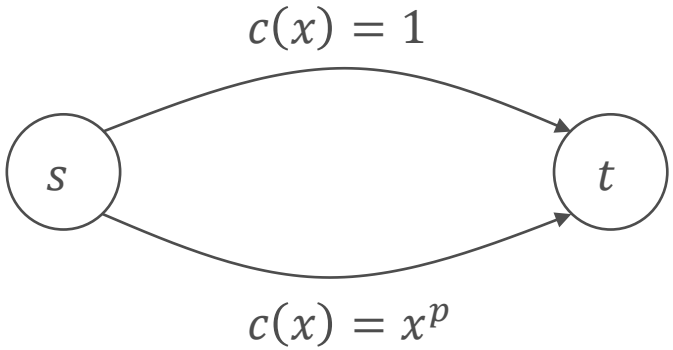
- Total commute time $= 2$!
- Braess's Paradox: the intuitively "helpful" action of adding a new zero-cost link can negatively impact all of the traffic!

PRICE OF ANARCHY

- Price of Anarchy (PoA): the ratio between the worst case "selfish" outcome and the best "centralized" outcome
- PoA in the previous example: $2/1.5 = 4/3$
- Can we bound it in general?

PIGOU'S EXAMPLE





- In equilibrium, everyone picks the bottom path, no matter what p is
  - Total cost $= 1$
- Having $(1 - \epsilon)$ on the traffic on the bottom link gives $(1 - \epsilon)^p \to 0$ cost for them
  - Total cost is close to ϵ
- PoA is unbounded!

PRICE OF ANARCHY IN ROUTING GAMES

- So far:
  - Linear costs gave small PoA
  - Non-linear costs gave large PoA
- Best case scenario: “In every routing game, if the edge costs are linear, then PoA is small”
- Let’s define a routing game first!

MODEL

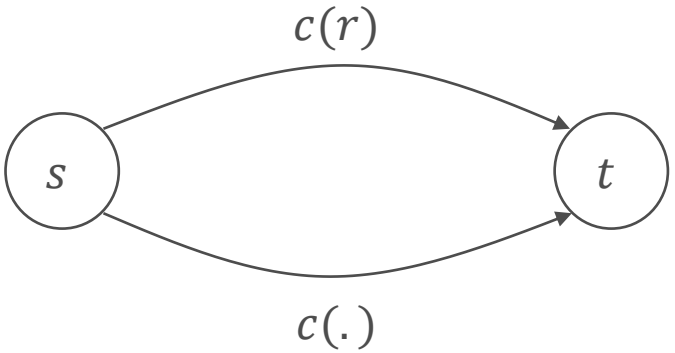
- Directed graph $G = (V, E)$ with a source s and a sink t
- There are r units of flow that we must push from s to t
- Each edge $e \in E$ has a non-negative and non-decreasing (continuous) cost function $c_e(.)$, which indicates the cost (per unit of flow) as a function of the flow that uses this edge
  - No capacities on the edges

ROUTING GAMES

- Our PoA bounds will be parameterized by families of cost functions $\mathcal{C}$!
  - E.g. $\mathcal{C} = \{c(x) = ax + b : a, b \geq 0\}$
- Informal Theorem [Roughgarden 2003]: Among all networks with cost functions in a set $\mathcal{C}$, the larget PoA is achieved in a Pigou-like network.
- Given this theorem, calculating the PoA for a given family of cost functions is super easy
  - Don't need to consider all network structures!

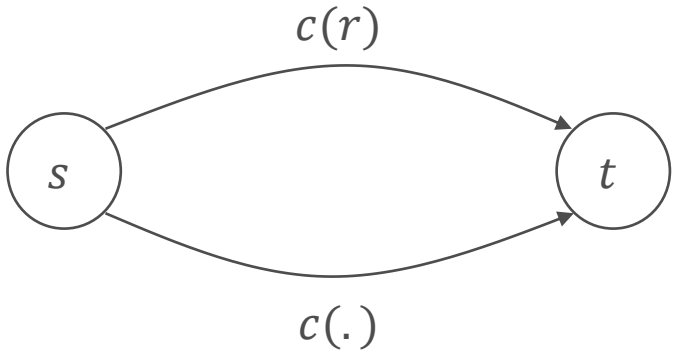
TOWARDS A PROOF

TOWARDS A ~~PROOF~~ STATEMENT



- A Pigou-like network has
  - Two vertices s, t
  - Two edges from s to t
  - A traffic rate $r > 0$
  - A cost function $c(.)$ on one edge, and a cost function $c(r)$ on the other edge
- Two free parameters: r and the cost function c

TOWARDS A ~~PROOF~~ STATEMENT



- By construction, equilibrium has cost

$$r \cdot c(r)$$

- Social optimum is

$$\min_{0 \le x \le r} \{x \cdot c(x) + (r - x) \cdot c(r)\}$$

- Pigou bound

$$\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \sup_{r \ge 0} \sup_{x \ge 0} \left\{ \frac{r \cdot c(r)}{x \cdot c(x) + (r - x) \cdot c(r)} \right\}$$

A STATEMENT

- Theorem [Roughgarden 2003]: For every set of cost functions $\mathcal{C}$ and every selfish routing game, the PoA is at most $\alpha(\mathcal{C})$

NETWORK FLOW PRELIMS

- $f_e = \sum_{P \in \mathcal{P}: e \in P} f_P$ is the amount of flow that goes through edge e
- A flow is in equilibrium iff traffic travels on shortest paths (wrt the flow)
  - $f_{\hat{P}} > 0$ only if $\hat{P} \in argmin_{P \in \mathcal{P}} \{\sum_{e \in P} c_e(f_e)\}$
- Objective: total travel time
  - $C(f) = \sum_{P \in \mathcal{P}} f_P \cdot c_P(f)$
    - $c_P(f) = \sum_{e \in P} c_e(f_e)$
  - $C(f) = \sum_{e \in E} f_e \cdot c_e(f_e)$
- Non-trivial fact: All equilibrium flows have the same cost

PROOF OF THEOREM

- Let f be an equilibrium flow and f^* be an optimal flow
- TODO #1: If we "freeze" edge costs at $c_e(f_e)$ then f is optimal
  - All paths $\hat{P}$ used in f have a common cost $c_{\hat{P}}(f) = L$ (since f is in equilibrium)
  - $c_P(f) \geq L$ for all $P \in \mathcal{P}$
  - $\sum_{P \in \mathcal{P}} f_P \cdot c_P(f) = r \cdot L$
  - $\sum_{P \in \mathcal{P}} f_P^* \cdot c_P(f) \geq r \cdot L$
  - So, $\sum_{e \in E} (f_e^* - f_e) \cdot c_e(f_e) \geq 0$ (1)

PROOF OF THEOREM

- TODO #2: Show that edge by edge, gap between f^* and f is at most Pigou bound
  - $\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \sup_{r \geq 0} \sup_{x \geq 0} \left\{ \frac{r \cdot c(r)}{x \cdot c(x) + (r - x) \cdot c(r)} \right\}$
  - For all $e \in E$ we have (replace r, x, c with f_e, f_e^*, c_e):

$$\alpha(\mathcal{C}) \geq \frac{f_e \cdot c_e(f_e)}{f_e^* \cdot c_e(f_e^*) + (f_e - f_e^*) c_e(f_e)}$$

  - Re-arrange:

$$f_e^* \cdot c_e(f_e^*) \geq \frac{1}{\alpha(\mathcal{C})} \cdot f_e \cdot c_e(f_e) + (f_e^* - f_e) \cdot c_e(f_e)$$

PROOF OF THEOREM

- TODO #2: Show that edge by edge, gap between f^* and f is at most Pigou bound

$$f_e^* \cdot c_e(f_e^*) \geq \frac{1}{\alpha(\mathcal{C})} \cdot f_e \cdot c_e(f_e) + (f_e^* - f_e) \cdot c_e(f_e)$$

  - Sum up over all $e \in E$

$$C(f^*) \geq \frac{1}{\alpha(\mathcal{C})} \cdot C(f) + \sum_{e \in E} (f_e^* - f_e) \cdot c_e(f_e)$$

$$\geq^{Ineq.\,(1)} \frac{1}{\alpha(\mathcal{C})} \cdot C(f)$$

  - Thus $\frac{C(f)}{C(f^*)} \leq \alpha(\mathcal{C})$

RECAP

- Defined Price of Anarchy
  - Measures how inefficient an equilibrium can be
- Gave a tight bound on the Price of Anarchy for a large class of games (routing games)

REFERENCES

- Tim Roughgarden's lecture notes:
  - http://timroughgarden.org/f13/l/l11.pdf

DISCUSSION

- Is this theory “reasonable”?
  - Does NE make sense? Should “learning” replace everything?
  - What is the strongest solution concept you actually believe as a prediction of behavior (if any)?
- Should we model strategic behavior via equilibria? Or via dynamics?