CSC304
Algorithmic Game Theory & Mechanism Design

Nisarg Shah

Recap: PoA & PoS

- Price of Anarchy (PoA)

  "Worst NE vs optimum"

  $$\frac{\text{Max total reward}}{\text{Min total reward in any NE}}$$

  or

  $$\frac{\text{Max total cost in any NE}}{\text{Min total cost}}$$

- Price of Stability (PoS)

  "Best NE vs optimum"

  $$\frac{\text{Max total reward}}{\text{Max total reward in any NE}}$$

  or

  $$\frac{\text{Min total cost in any NE}}{\text{Min total cost}}$$

$$\boxed{\text{PoA} \geq \text{PoS} \geq 1}$$

Recap: Cost Sharing Game

- n players on directed weighted graph G

- Player i
  - Wants to go from s_i to t_i
  - Strategy set S_i = {directed $s_i \to t_i$ paths}
  - Denote his chosen path by $P_i \in S_i$

- Each edge e has cost c_e (weight)
  - Cost is split among all players taking edge e
  - That is, among all players i with $e \in P_i$

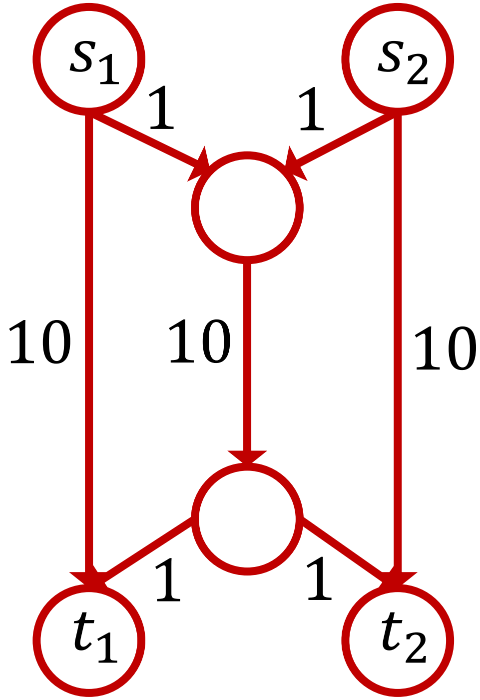

Recap: Cost Sharing Game

- Given strategy profile $\vec{P}$, cost $c_i(\vec{P})$ to player i is sum of his costs for edges $e \in P_i$
- Social cost $C(\vec{P}) = \sum_i c_i\,(\vec{P})$
- Note: $C(\vec{P}) = \sum_{e \in E(\vec{P})} c_e$, where…
  - $E(\vec{P})$={edges taken in $\vec{P}$ by at least one player}
  - Why?

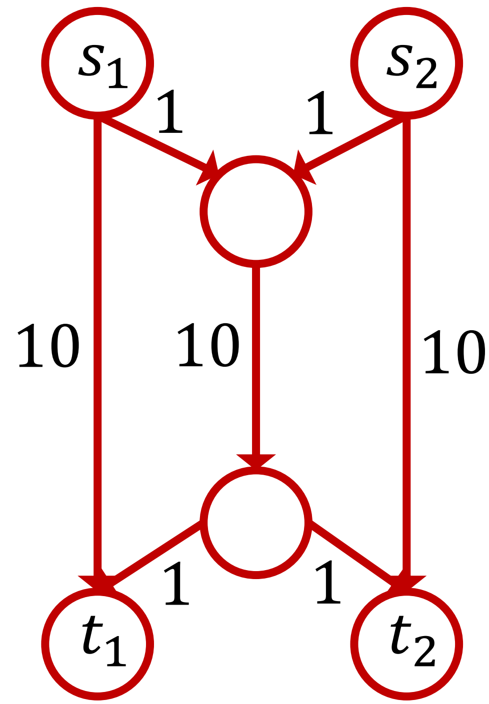

Recap: PoA of Cost-Sharing

- For cost-sharing games, we'll be looking at PoA and PoS with respect to pure Nash equilibria.
- Theorem:
  - Every cost-sharing game has $\mathrm{PoA} \leq n$.
  - There exists a cost-sharing game with $\mathrm{PoA} = n$

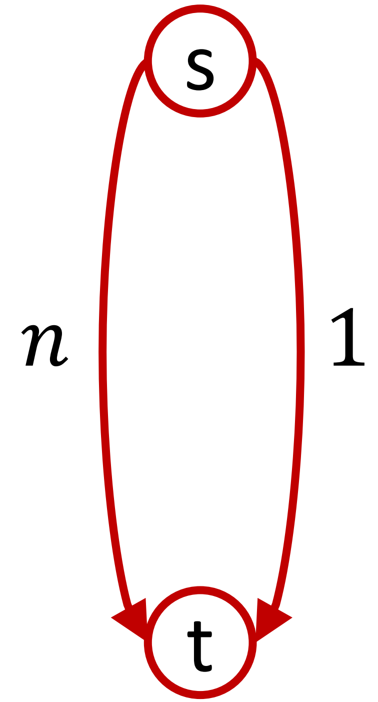

- Before looking at PoS…
  - Want to argue that every cost-sharing game admits a pure Nash equilibrium via "potential" argument
    - Will prove that $PoS = O(\log n)$
    - Tightness established in tutorial 3

Good News

- Theorem: Every cost-sharing game has a pure Nash equilibrium.

- Proof:
  - Via "potential function" argument

Step 1: Define Potential Fn

- Potential function: $\Phi : \prod_i S_i \to \mathbb{R}_+$
  - This is a function such that for every pure strategy profile $\vec{P} = (P_1, \dots, P_n)$, player i, and strategy P_i' of i,

$$c_i(P_i', \vec{P}_{-i}) - c_i(\vec{P}) = \Phi(P_i', \vec{P}_{-i}) - \Phi(\vec{P})$$

  - When a single player i changes her strategy, the change in potential function equals the change in cost to i!
  - Note: In contrast, the change in the social cost C equals the total change in cost to all players.
    - Hence, the social cost will often not be a valid potential function.

Step 2: Potential $\mathrm{F^n} \rightarrow$ pure Nash Eq

- A potential function exists ⇒ a pure NE exists.
  - Consider a $\vec{P}$ that minimizes the potential function.
  - If player i deviates to playing P_i', then by the definition of the potential function:
    $$c_i(P_i', \vec{P}_{-i}) - c_i(P_i, \vec{P}_{-i}) = \Phi(P_i', \vec{P}_{-i}) - \Phi(P_i, \vec{P}_{-i}) \geq 0$$
  - The inequality is because $\Phi(P_i, \vec{P}_{-i})$ is the lowest possible.
  - Hence, player i's cost cannot decrease by deviating.

- Hence, every pure strategy profile minimizing the potential function is a pure Nash equilibrium.

Step 3: Potential F^n for Cost-Sharing

- Recall: $E(\vec{P}) = \{\text{edges taken in } \vec{P} \text{ by at least one player}\}$
- Let $n_e(\vec{P})$ be the number of players taking e in $\vec{P}$

$$\Phi(\vec{P}) = \sum_{e \in E(\vec{P})} \sum_{k=1}^{n_e(\vec{P})} \frac{c_e}{k}$$

- Note: The cost of edge e to each player taking e is $c_e/n_e(\vec{P})$. But the potential function includes all fractions: $c_e/1, c_e/2, \ldots, c_e/n_e(\vec{P})$.

Step 3: Potential F^n for Cost-Sharing

$$\Phi(\vec{P}) = \sum_{e \in E(\vec{P})} \sum_{k=1}^{n_e(\vec{P})} \frac{c_e}{k}$$

- Why is this a potential function?
  - If a player changes path, he pays $\frac{c_e}{n_e(\vec{P})+1}$ for each new edge e, gets back $\frac{c_f}{n_f(\vec{P})}$ for each old edge f.
  - This is precisely the change in the potential function too.
  - So $\Delta c_i = \Delta\Phi$.

■

Potential Minimizing Eq.

- Minimizing the potential function gives some pure Nash equilibrium
  - Is this equilibrium special? Yes!

- Recall that the price of anarchy can be up to n.
  - That is, the worst Nash equilibrium can be up to n times worse than the social optimum.

- A potential-minimizing pure Nash equilibrium is better!

Potential Minimizing Eq.

- $$\sum_{e \in E(\vec{P})} c_e \leq \Phi(\vec{P}) = \sum_{e \in E(\vec{P})} \sum_{k=1}^{n_e(\vec{P})} \frac{c_e}{k} \leq \sum_{e \in E(\vec{P})} c_e * \sum_{k=1}^{n} \frac{1}{k}$$

Social cost

Harmonic function $H(n)$ $= \sum_{k=1}^{n} 1/n = O(\log n)$

- $$\forall \vec{P}, \;\; C(\vec{P}) \leq \Phi(\vec{P}) \leq C(\vec{P}) * H(n)$$

- $$C(\vec{P}^*) \leq \Phi(\vec{P}^*) \leq \Phi(OPT) \leq C(OPT) * H(n)$$

Potential minimizing eq.

Social optimum

Potential Minimizing Eq.

- Potential-minimizing PNE is $O(\log n)$-approximation to the social optimum.

- Thus, in every cost-sharing game, the price of stability is $O(\log n)$.
  - Compare to the price of anarchy, which can be n

Congestion Games

- Generalize cost sharing games
- n players, m resources (e.g., edges)
- Each player i chooses a set of resources P_i (e.g., $s_i \rightarrow t_i$ paths)
- When n_j player use resource j, each of them get a cost $f_j(n_j)$
- Cost to player is the sum of costs of resources used

Congestion Games

- Theorem [Rosenthal 1973]: Every congestion game is a potential game.

- Potential function:

$$\Phi(\vec{P}) = \sum_{j \in E(\vec{P})} \sum_{k=1}^{n_j(\vec{P})} f_j(k)$$

- Theorem [Monderer and Shapley 1996]: Every potential game is equivalent to a congestion game.

The Braess' Paradox

- In cost sharing, f_j is decreasing
  - The more people use a resource, the less the cost to each.
- f_j can also be increasing
  - Road network, each player going from home to work
  - Uses a sequence of roads
  - The more people on a road, the greater the congestion, the greater the delay (cost)
- Can lead to unintuitive phenomena

The Braess' Paradox

- Parkes-Seuken Example
  - 2000 players want to go from 1 to 4
  - 1 → 2 and 3 → 4 are "congestible" roads
  - 1 → 3 and 2 → 4 are "constant delay" roads

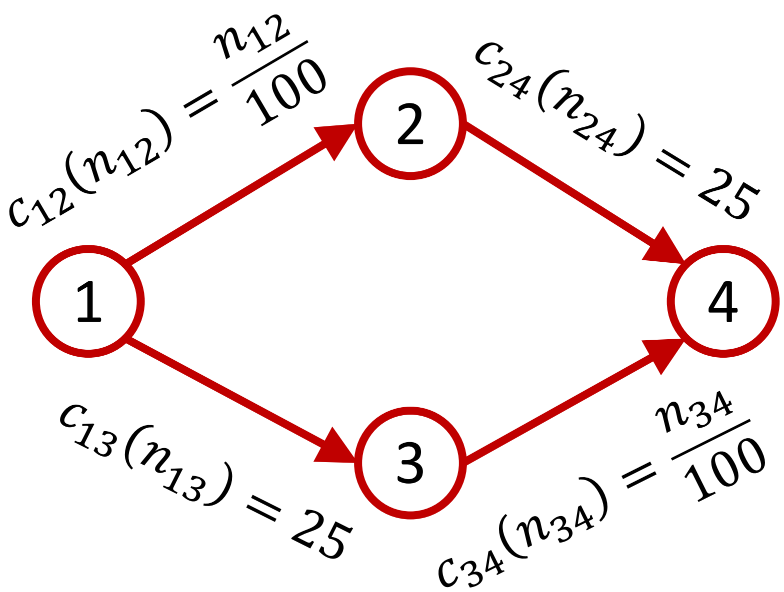

The Braess' Paradox

- Pure Nash equilibrium?
  - 1000 take $1 \to 2 \to 4$, 1000 take $1 \to 3 \to 4$
  - Each player has cost $10 + 25 = 35$
  - Anyone switching to the other creates a greater congestion on it, and faces a higher cost

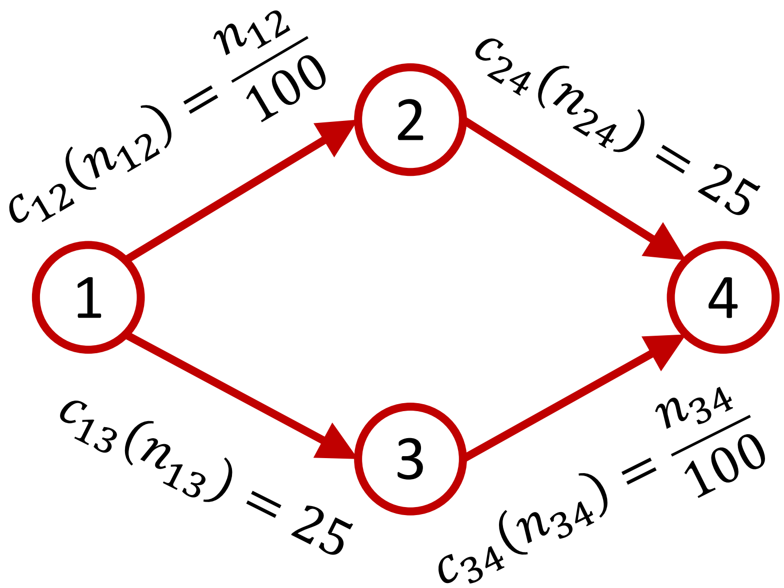

The Braess' Paradox

- What if we add a zero-cost connection $2 \rightarrow 3$?
  - Intuitively, adding more roads should only be helpful
  - In reality, it leads to a greater delay for everyone in the unique equilibrium!

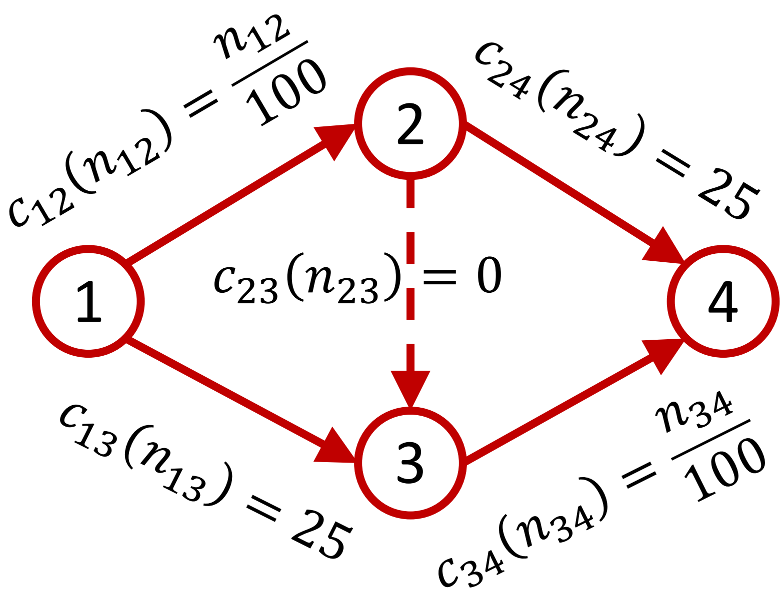

The Braess' Paradox

- Nobody chooses $1 \rightarrow 3$ as $1 \rightarrow 2 \rightarrow 3$ is better irrespective of how many other players take it
- Similarly, nobody chooses $2 \rightarrow 4$
- Everyone takes $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$, faces delay = 40!

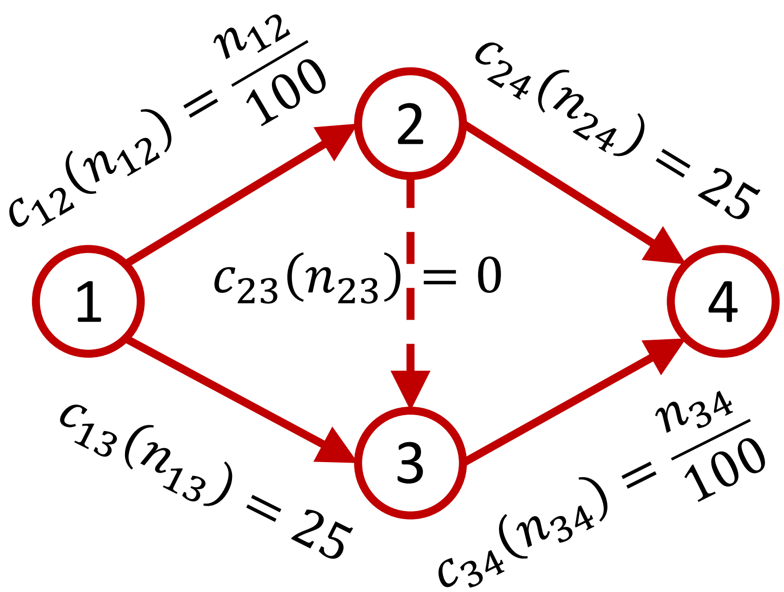

The Braess' Paradox

- In fact, what we showed is:
  - In the new game, $1 \to 2 \to 3 \to 4$ is a strictly dominant strategy for each player!

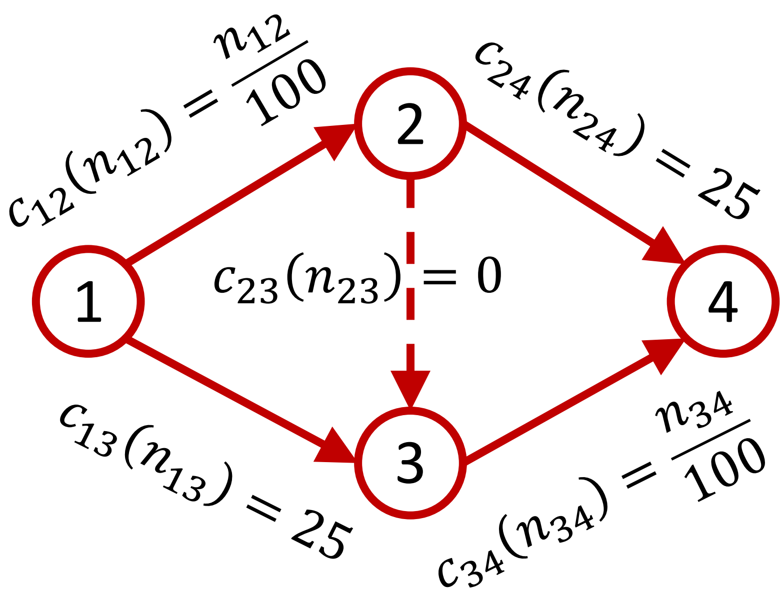

Zero-Sum Games

Zero-Sum Games

- Special case of games
  - Total reward to all players is constant in every outcome
  - Without loss of generality, sum of rewards = 0
    - Remember: rewards to each player can be additively shifted without changing the structure of the game
  - Inspired terms like “zero-sum thinking” and “zero-sum situation”

- Focus on two-player zero-sum games (2p-zs)
  - “The more I win, the more you lose”

Examples

Zero-sum game: Rock-Paper-Scissor

P2 / P1	Rock	Paper	Scissor
Rock	(0 , 0)	(-1 , 1)	(1 , -1)
Paper	(1 , -1)	(0 , 0)	(-1 , 1)
Scissor	(-1 , 1)	(1 , -1)	(0 , 0)

Non-zero-sum game: Prisoner's dilemma

John / Sam	Stay Silent	Betray
Stay Silent	(-1 , -1)	(-3 , 0)
Betray	(0 , -3)	(-2 , -2)

Importance

- Why are they interesting?
  - Many physical games we play are zero-sum: chess, tic-tac-toe, rock-paper-scissor, …
  - (win, lose), (lose, win), (draw, draw)
  - (1, -1), (-1, 1), (0, 0)

- Why are they technically interesting?
  - We'll see.

Zero-Sum Games

- Reward for P2 = - Reward for P1
  - Only need to write a single entry in each cell (say reward of P1)
    - We get a matrix A
  - Row player wants to maximize the value, column player wants to minimize it

P1 \ P2	Rock	Paper	Scissor
Rock	**0**	**-1**	**1**
Paper	**1**	**0**	**-1**
Scissor	**-1**	**1**	**0**

Rewards in Matrix Form

- Reward to the row player when…
  - The row player uses mixed strategy $x_1 = (x_{1,1}, x_{1,2}, \dots)$
  - The column player uses mixed strategy $x_2 = (x_{2,1}, x_{2,2}, \dots)$
  - Given by

$$x_1^T \, A \, x_2 = \sum_{i,j} x_{1,i} * x_{2,j} * A_{i,j}$$

    - With probability $x_{1,i} * x_{2,j}$, the row player chooses action i and the column player chooses action j, giving the row player reward $A_{i,j}$

- Reward to the column player is $-\, x_1^T \, A \, x_2$

How would the two players act in this zero-sum game?

John von Neumann, 1928

Maximin Strategy

- Worst-case thinking by the row player…
  - If I choose mixed strategy x_1, *in the worst case* the column player might end up choosing x_2 that minimizes my reward
  - Let me choose x_1 that maximizes this "worst-case reward":

$$V_1^* = \max_{x_1} \min_{x_2} x_1^T A x_2$$

  - V_1^* = maximin value of the row player
  - x_1^* (maximizer) = maximin strategy of the row player
  - Row player guarantee:
    - "By playing x_1^*, I guarantee myself reward at least V_1^*"

Maximin Strategy

- Similar worst-case thinking by the column player…
  - If I choose mixed strategy x_2, *in the worst case* the row player ends up choosing x_1 minimizing my reward (i.e., maximizing her reward)
  - Let me choose x_2 that optimizes this "worst-case":

$$V_2^* = \min_{x_2} \max_{x_1} x_1^T A\, x_2$$

  - V_2^* = minimax value of the column player
  - x_2^* (maximizer) = minimax strategy of the column player
  - Column player guarantee:
    - "By playing x_2^*, I guarantee that the row player gets at most V_2^*"

Maximin vs Minimax

Row player

If I play x_1^*, I get reward at least V_1^*

$$V_1^* = \max_{x_1} \min_{x_2} x_1^T * A * x_2$$

x_1^*

Column player

If I play x_2^*, the row player gets reward at most V_2^*

$$V_2^* = \min_{x_2} \max_{x_1} x_1^T * A * x_2$$

x_2^*

Claim: It is easy to see that $V_1^* \leq V_2^*$ **(Why?)**

Maximin vs Minimax

$$V_1^* = \max_{x_1} \; \min_{x_2} \; x_1^T * A * x_2 \qquad V_2^* = \min_{x_2} \; \max_{x_1} \; x_1^T * A * x_2$$

$$x_1^* \qquad\qquad x_2^*$$

- Another way to see this:

$$V_1^* = \min_{x_2} \; (x_1^*)^T * A * x_2 \; \leq \; (x_1^*)^T * A * x_2^*$$

$$\leq \; \max_{x_1} \; x_1^T * A * x_2^* = \; V_2^*$$

The Minimax Theorem

- Jon von Neumann [1928]

- Theorem: For any two-player zero-sum game,
  - $V_1^* = V_2^* = V^*$ (called the minimax value of the game)
  - Set of Nash equilibria =

    $\{(\mathrm{x}_1^*, \mathrm{x}_2^*)$: where…

    x_1^* = maximin for row player, x_2^* = minimax for column player$\}$

- Corollary: x_1^* is best response to x_2^* and vice-versa.

Commitment Interpretation

- Commitment
  - x_1^* is the strategy that the row player would choose if she were to commit to her strategy *first,* and the column player were to choose his strategy after observing the row player's strategy
  - Similarly, x_2^* is the strategy that the column player would choose if he were to commit to his strategy *first,* and the row player were to choose her strategy after observing the column player's strategy

- Minimax theorem:
  - x_1^* and x_2^* are best responses to each other, so in two-player zero-sum games, it doesn't matter if one player commits first or if both play simultaneously

The Minimax Theorem

- Jon von Neumann [1928]

  "As far as I can see, there could be no theory of games … without that theorem …

  I thought there was nothing worth publishing until the Minimax Theorem was proved"

Computing Nash Equilibria

- Recall that in general games, computing a Nash equilibrium is hard even with two players.

- For two-player zero-sum games, a Nash equilibrium can be computed in polynomial time.
  - Polynomial in #actions of the two players: m_1 and m_2
  - Exploits the fact that Nash equilibrium is simply composed of maximin strategies, which can be computed using linear programming

Computing Nash Equilibria

Maximize v

Subject to

$$(x_1^T A)_j \geq v,\ j \in \{1, \dots, m_2\}$$

$$x_1(1) + \cdots + x_1(m_1) = 1$$

$$x_1(i) \geq 0, i \in \{1, \dots, m_1\}$$

Minimax Theorem in Real Life?

Kicker \ Goalie	L	R
L	**0.58**	**0.95**
R	**0.93**	**0.70**

Kicker

Maximize v

Subject to

$0.58p_L + 0.93p_R \geq v$

$0.95p_L + 0.70p_R \geq v$

$p_L + p_R = 1$

$p_L \geq 0, p_R \geq 0$

Goalie

Minimize v

Subject to

$0.58q_L + 0.95q_R \leq v$

$0.93q_L + 0.70q_R \leq v$

$q_L + q_R = 1$

$q_L \geq 0, q_R \geq 0$

Minimax Theorem in Real Life?

Kicker \ Goalie	L	R
L	0.58	0.95
R	0.93	0.70

Kicker
Maximin:
$p_L = 0.38, p_R = 0.62$

Reality:
$p_L = 0.40, p_R = 0.60$

Goalie
Maximin:
$q_L = 0.42, q_R = 0.58$

Reality:
$p_L = 0.423, q_R = 0.577$

Some evidence that people may play minimax strategies.

Stackelberg Games

Sequential Move Games

- Focus on two players: “leader” and “follower”

1. Leader commits to a (possibly mixed) strategy x_1
   - Cannot change later
2. Follower learns about x_1
   - Follower must believe that leader’s commitment is credible
3. Follower chooses the best response x_2
   - Can assume to be a pure strategy without loss of generality
   - If multiple actions are best response, break ties in favor of the leader

Sequential Move Games

- Wait. Does this give us anything new?
  - Can't I, as player 1, commit to playing x_1 in a simultaneous-move game too?
  - Player 2 wouldn't believe you.

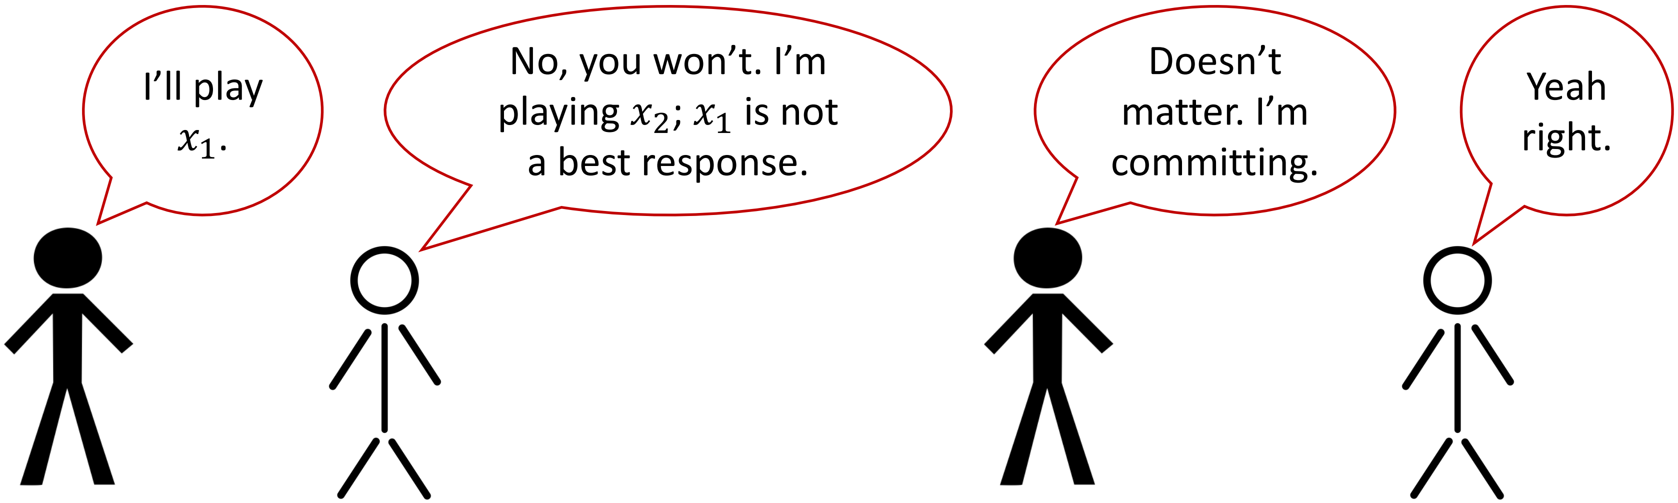

That's unless…

- You're as convincing as this guy.

How to represent the game?

- Extensive form representation
  - Can also represent “information sets”, multiple moves, …

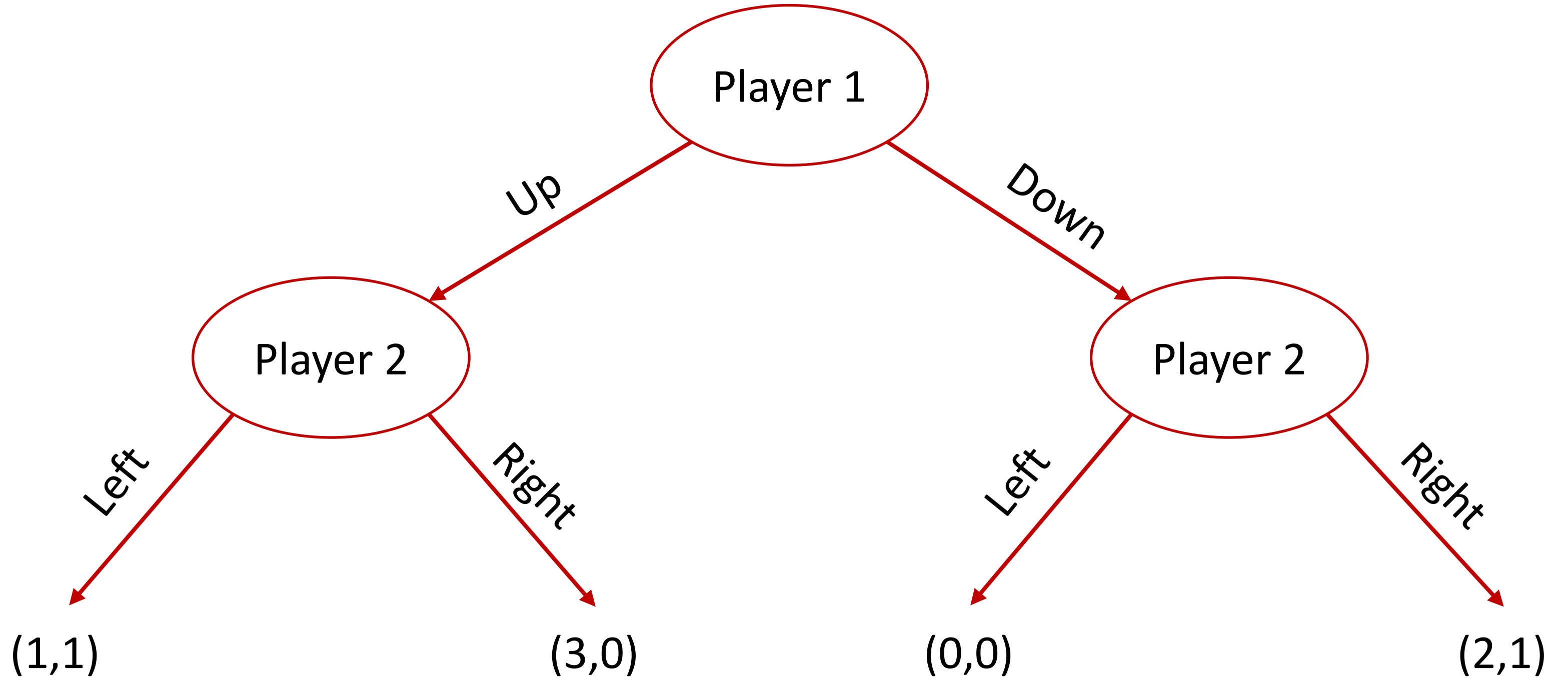

A Curious Case

P1 \ P2	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Q: What are the Nash equilibria of this game?
- Q: You are P1. What is your reward in Nash equilibrium?

A Curious Case

P1 \ P2	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Say that as P1, you have the ability to commit to a pure strategy.

- Q: Which pure strategy would you commit to? And what would your reward be now?

Commitment Advantage

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Reward in the unique Nash equilibrium = 1
- (Higher) reward when committing to Down = 2

Commitment Advantage

P2 / P1	Left	Right
Up	(1 , 1)	(3 , 0)
Down	(0 , 0)	(2 , 1)

- Even higher reward in committing to a mixed strategy
  - P1 commits to: Up w.p. $0.5 - \epsilon$, Down w.p. $0.5 + \epsilon$
  - P2 is still better off playing Right
  - $\mathbb{E}$[Reward] to P1 $\rightarrow 2.5$
  - Note: If P1 plays both actions with probability exactly 0.5, we assume P2 plays Right (break ties in favor of leader)