CSC2556

Lecture 11

Game Theory II:
Prices of Anarchy and Stability,
Zero-Sum Games

Prices of Anarchy and Stability

Price of Anarchy and Stability

- If players play a Nash equilibrium instead of "socially optimum", how bad can it be?

- Objective function: sum of utilities/costs

- Price of Anarchy (PoA): compare the optimum to the worst Nash equilibrium

- Price of Stability (PoS): compare the optimum to the best Nash equilibrium

Price of Anarchy and Stability

- Price of Anarchy (PoA)

$$\frac{\text{Max social utility}}{\text{Min social utility in any NE}}$$

Costs $\rightarrow$ flip:
Nash equilibrium
divided by optimum

- Price of Stability (PoS)

$$\frac{\text{Max social utility}}{\text{Max social utility in any NE}}$$

Revisiting Stag-Hunt

Hunter 2 \ Hunter 1	Stag	Hare
Stag	(4 , 4)	(0 , 2)
Hare	(2 , 0)	(1 , 1)

- Optimum social utility = 4+4 = 8
- Three equilibria:
  - (Stag, Stag) : Social utility = 8
  - (Hare, Hare) : Social utility = 2
  - (Stag:1/3 - Hare:2/3, Stag:1/3 - Hare:2/3)
    - Social utility = (1/3)*(1/3)*8 + (1-(1/3)*(1/3))*2 = Btw 2 and 8
- Price of stability? Price of anarchy?

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 $s_i \rightarrow t_i$ 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 $e \in P_i$

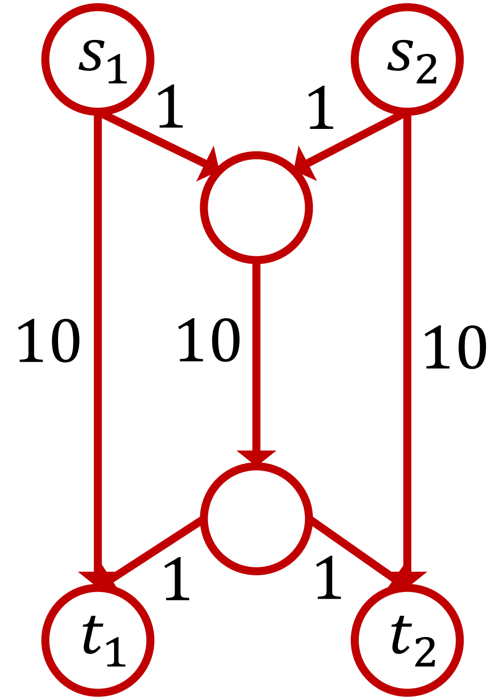

Cost Sharing Game

- Given strategy profile $\vec{P}$, cost $c_i\left(\vec{P}\right)$ to player i is sum of his costs for edges $e \in P_i$
- Social cost $C\left(\vec{P}\right) = \sum_i c_i\left(\vec{P}\right)$
  - Note that $C\left(\vec{P}\right) = \sum_{e \in E(\vec{P})} c_e$, where $E(\vec{P})$={edges taken in $\vec{P}$ by at least one player}
- In the example on the right:
  - What if both players take the direct paths?
  - What if both take the middle paths?
  - What if only one player takes the middle path while the other takes the direct path?

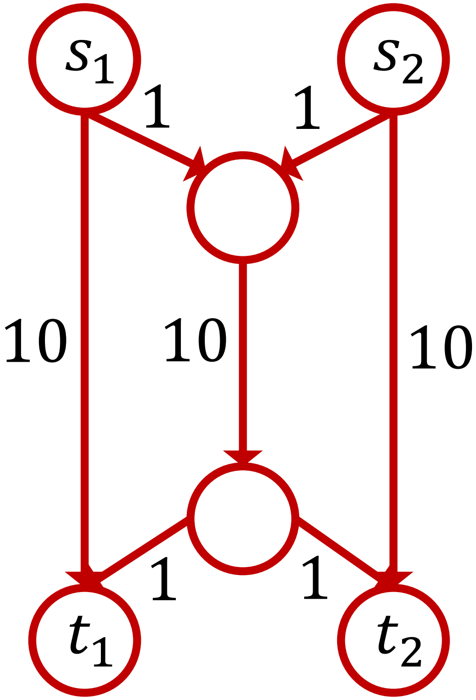

Cost Sharing: Simple Example

- Example on the right: n players
- Two pure NE
  - All taking the n-edge: social cost = n
  - All taking the 1-edge: social cost = 1
    - Also the social optimum
- In this game, price of anarchy $\geq n$
- We can show that for all cost sharing games, price of anarchy $\leq n$

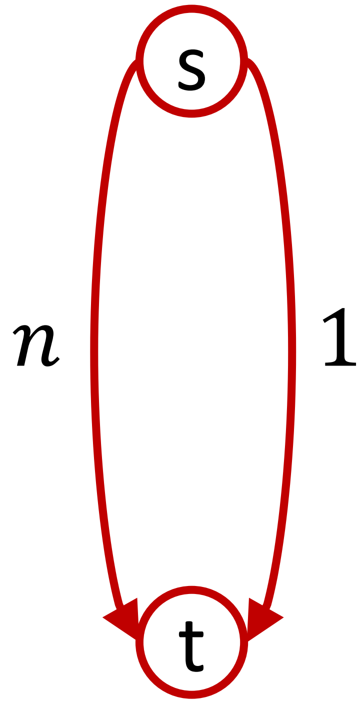

Cost Sharing: PoA

- **Theorem:** The price of anarchy of a cost sharing game is at most n.
- **Proof:**
  - Suppose the social optimum is $(P_1^*, P_2^*, \dots, P_n^*)$, in which the cost to player i is c_i^*.
  - Take any NE with cost c_i to player i.
  - Let c_i' be his cost if he switches to P_i^*.
  - NE $\Rightarrow c_i' \geq c_i$ (Why?)
  - But : $c_i' \leq n \cdot c_i^*$ (Why?)
  - $c_i \leq n \cdot c_i^*$ for each i $\Rightarrow$ no worse than $n \times$ optimum ■

Cost Sharing

- Price of anarchy
  - All cost-sharing games: PoA $\leq n$
  - $\exists$ example where PoA $= n$
- Price of stability? Later…
- Both examples we saw had pure Nash equilibria
  - What about more complex games, like the one on the right?

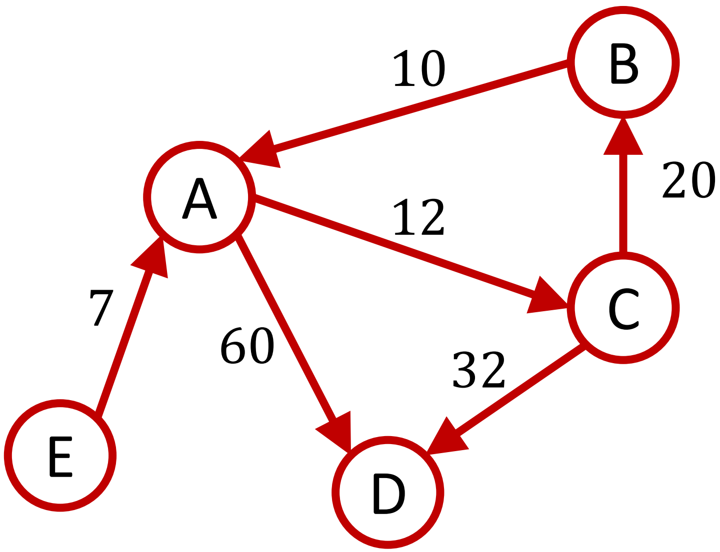

10 players: $E \to C$
27 players: $B \to D$
19 players: $C \to D$

Good News

- Theorem: All cost sharing games admit a pure Nash equilibrium.

- Proof:
  - Via a "potential function" argument.

Step 1: Define Potential Fn

- Potential function: $\Phi : \prod_i S_i \to \mathbb{R}_+$
  - For all pure strategy profiles $\vec{P} = (P_1, \dots, P_n) \in \prod_i S_i$, …
  - all players i, and …
  - all alternative strategies $P_i' \in S_i$ for player i…

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

- When a single player changes his strategy, the change in *his* cost is equal to the change in the potential function
  - Do not care about the changes in the costs to others

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

- All games that admit a potential function have a pure Nash equilibrium. Why?
  - Think about $\vec{P}$ that minimizes the potential function.
  - What happens when a player deviates?
    - If his cost decreases, the potential function value must also decrease.
    - $\vec{P}$ already minimizes the potential function value.

- Pure strategy profile minimizing 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 $\vec{P}$

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

- 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: $c_e/1, c_e/2, \ldots, c_e/n_e\left(\vec{P}\right)$.

Step 3: Potential $\mathrm{F^n}$ for Cost-Sharing

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

- 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.

- There could be multiple pure Nash equilibria
  - Pure Nash equilibria are "local minima" of the potential function.
  - *A single player* deviating should not decrease the function value.

- Is the *global minimum* of the potential function a special pure Nash equilibrium?

Potential Minimizing Eq.

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

Social cost

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

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

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

Potential minimizing eq.

Social optimum

Potential Minimizing Eq.

- Potential minimizing equilibrium gives $O(\log n)$ approximation to the social optimum
  - Price of stability is $O(\log n)$
    - $\exists$ example where price of stability is $\Theta(\log n)$
  - 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:

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

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

Potential Functions

- Potential functions are useful for deriving various results
  - E.g., used for analyzing amortized complexity of algorithms

- Bad news: Finding a potential function that works may be hard.

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

- Due to Parkes and Seuken:
  - 2000 players want to go from 1 to 4
  - $1 \rightarrow 2$ and $3 \rightarrow 4$ are "congestible" roads
  - $1 \rightarrow 3$ and $2 \rightarrow 4$ are "constant delay" roads

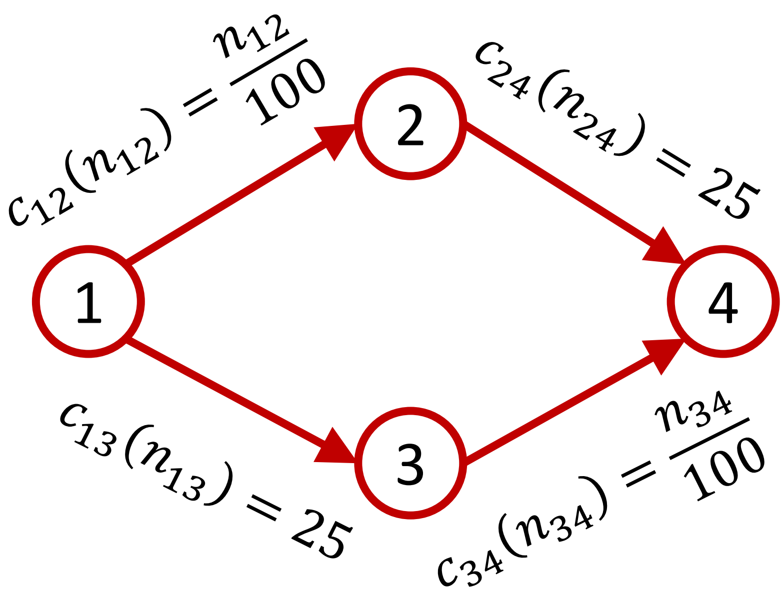

The Braess' Paradox

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

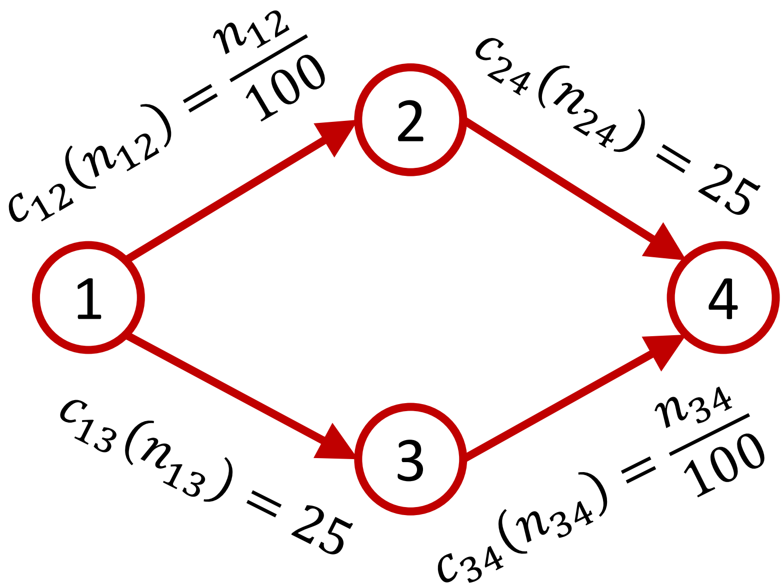

The Braess' Paradox

- 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!

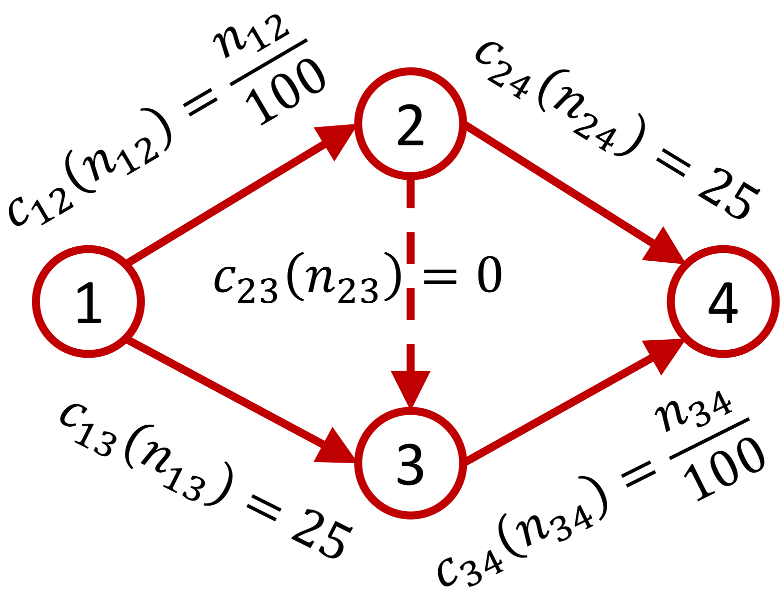

The Braess' Paradox

- 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 $= 40$!

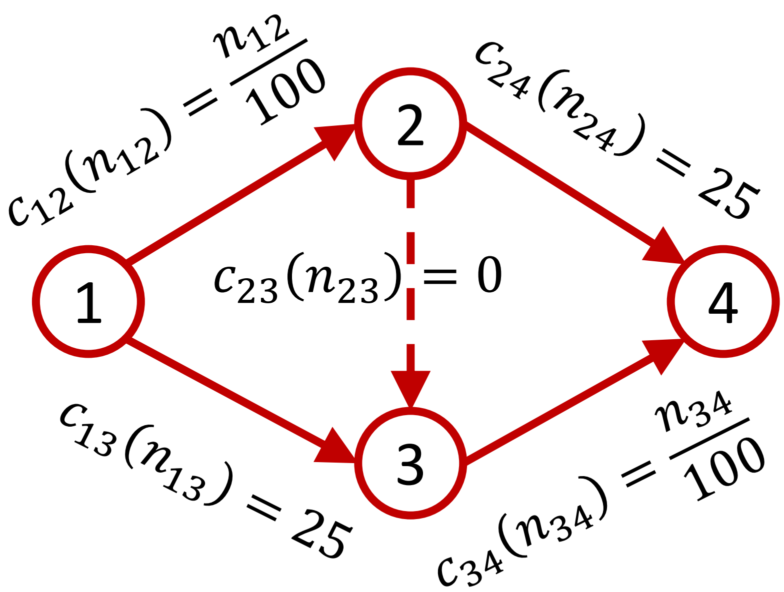

The Braess' Paradox

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

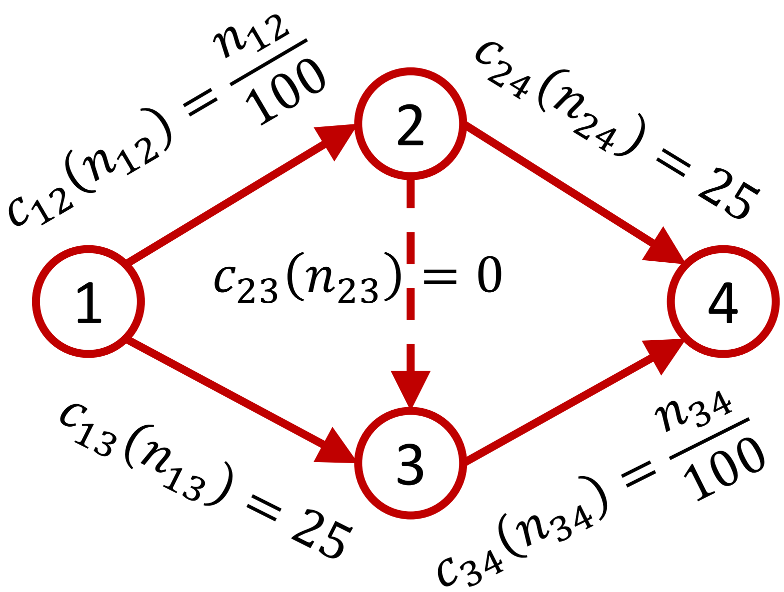

Zero-Sum Games

Zero-Sum Games

- Total reward is constant in all outcomes (w.l.o.g. 0)
- Focus on two-player zero-sum games (2p-zs)
  - "The more I win, the more you lose"
  - Chess, tic-tac-toe, 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)

Zero-Sum Games

- Reward for P2 = - Reward for P1
  - Only need a single matrix A : reward for P1
  - P1 wants to maximize, P2 wants to minimize

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

Rewards in Matrix Form

- Reward for P1 when…
  - P1 uses mixed strategy x_1
  - P2 uses mixed strategy x_2
  - $x_1^T A x_2$ (where x_1 and x_2 are column vectors)

Maximin/Minimax Strategy

- Worst-case thinking by P1…
  - If I commit to x_1 first, P2 would choose x_2 to minimize my reward (i.e., maximize his reward)

- P1's best worst-case guarantee:

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

  - A maximizer x_1^* is a maximin strategy for P1

Maximin/Minimax Strategy

- P1's best worst-case guarantee:

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

- P2's best worst-case guarantee:

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

  - P2's minimax strategy x_2^* minimizes this

- $V_1^* \leq V_2^*$ (both play their "safe" strategies together)

The Minimax Theorem

- Jon von Neumann [1928]

- Theorem: For any 2p-zs game,
  - $V_1^* = V_2^* = V^*$ (called the minimax value of the game)
  - Set of Nash equilibria =

    $\{(\mathrm{x}_1^*, \mathrm{x}_2^*) : \mathrm{x}_1^* = \text{maximin for P1}, \mathrm{x}_2^* = \text{minimax for P2}\}$

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

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"

- Indeed, much more compelling and predictive than Nash equilibria in general-sum games (which came much later).

Computing Nash Equilibria

- General-sum games: Computing a Nash equilibrium is PPAD-complete even with just two players.
  - Trivia: Another notable PPAD-complete problem is finding a three-colored point in Sperner's Lemma.

- 2p-zs games: Polynomial time using linear programming
  - Polynomial in #actions of the two players: m_1 and m_2

Computing Nash Equilibria

Maximize v

Subject to

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

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

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

Minimax Theorem in Real Life?

- If you were to play a 2-player zero-sum game (say, as player 1), would you always play a maximin strategy?
- What if you were convinced your opponent is an idiot?
- What if you start playing the maximin strategy, but observe that your opponent is not best responding?

Minimax Theorem in Real Life?

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:

$p_L = 0.423, q_R = 0.577$

Minimax Theorem

- Implies Yao's minimax principle
- Equivalent to linear programming duality

John von Neumann

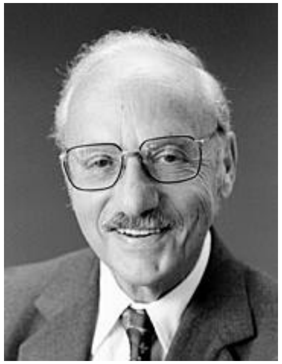

George Dantzig

von Neumann and Dantzig

George Dantzig loves to tell the story of his meeting with John von Neumann on October 3, 1947 at the Institute for Advanced Study at Princeton. Dantzig went to that meeting with the express purpose of describing the linear programming problem to von Neumann and asking him to suggest a computational procedure. He was actually looking for methods to benchmark the simplex method. Instead, he got a 90-minute lecture on Farkas Lemma and Duality (Dantzig's notes of this session formed the source of the modern perspective on linear programming duality). Not wanting Dantzig to be completely amazed, von Neumann admitted:

"I don't want you to think that I am pulling all this out of my sleeve like a magician. I have recently completed a book with Morgenstern on the theory of games. What I am doing is conjecturing that the two problems are equivalent. The theory that I am outlining is an analogue to the one we have developed for games.“

- (Chandru & Rao, 1999)