CSC2556

Lecture 9

Noncooperative Games 1:

Nash Equilibria, Price of Anarchy, Cost-Sharing Games

Game Theory

- How do rational, self-interested agents act in a given environment?
- Each agent has a set of possible actions
- Rules of the game:
  - Rewards for the agents as a function of the actions taken by all agents

- Noncooperative games
  - No external trusted agency, no legal agreements

Normal Form Games

- A set of players $\mathrm{N} = \{1, \dots, n\}$
- Each player i has an action set S_i, chooses $s_i \in S_i$
- $\mathcal{S} = S_1 \times \cdots \times S_n$.
- Action profile $\vec{s} = (s_1, \dots, s_n) \in \mathcal{S}$
- Each player i has a utility function $u_i : \mathcal{S} \to \mathbb{R}$
  - Given the action profile $\vec{s} = (s_1, \dots, s_n)$, each player i gets a reward $u_i(s_1, \dots, s_n)$

Normal Form Games

Prisoner's dilemma $S = \{\text{Silent}, \text{Betray}\}$

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

$\boxed{u_{Sam}(Betray, Silent)}$ $\boxed{u_{John}(Betray, Silent)}$

s_{Sam} s_{John}

Player Strategies

- Pure strategy
  - Deterministic choice of an action, e.g., “Betray”
- Mixed strategy
  - Randomized choice of an action, e.g., “Betray with probability 0.3, and stay silent with probability 0.7”

Dominant Strategies

- For player i, s_i dominates s_i' if s_i is "better than" s_i', *irrespective of other players' strategies*.

- Two variants: weak and strict domination
  - $u_i(s_i, \vec{s}_{-i}) \geq u_i(s_i', \vec{s}_{-i}), \forall \vec{s}_{-i}$
  - Strict inequality for some $\vec{s}_{-i}$ $\leftarrow$ Weak domination
  - Strict inequality for all $\vec{s}_{-i}$ $\leftarrow$ Strict domination

- s_i is a strictly (or weakly) dominant strategy for player i if it strictly (or weakly) dominates every other strategy

Dominant Strategies

- Q: How does this relate to strategyproofness?
- A: Strategyproofness means “truth-telling should be a weakly dominant strategy for every player”.

Example: Prisoner's Dilemma

- Recap:

John's Actions / Sam's Actions	Stay Silent	Betray
Stay Silent	**(-1 , -1)**	**(-3 , 0)**
Betray	**(0 , -3)**	**(-2 , -2)**

- Each player strictly wants to
  - Betray if the other player will stay silent
  - Betray if the other player will betray

- Betray = strictly dominant strategy for each player

Iterated Elimination

- What if there are no dominant strategies?
  - No single strategy dominates every other strategy
  - But some strategies might still be dominated

- Assuming everyone knows everyone is rational…
  - Can remove their dominated strategies
  - Might reveal a newly dominant strategy

- Eliminating only strictly dominated vs eliminating weakly dominated

Iterated Elimination

- Toy example:
  - Microsoft vs Startup
  - Enter the market or stay out?

Microsoft \ Startup	[redacted]	Stay Out
Enter	(2 , -2)	(4 , 0)
[redacted]	[redacted]	(0 , 0)

- Q: Is there a dominant strategy for startup?
- Q: Do you see a rational outcome of the game?

Iterated Elimination

- "Guess 2/3 of average"
  - Each student guesses a real number between 0 and 100 (inclusive)
  - The student whose number is the closest to 2/3 of the average of all numbers wins!

- Piazza Poll: What would you do?

Nash Equilibrium

- If we find dominant strategies, or a unique outcome after iteratively eliminating dominated strategies, it *may* be considered the rational outcome of the game.
- What if this is not the case?

Students \ Professor	Attend	Be Absent
Attend	**(3 , 1)**	**(-1 , -3)**
Be Absent	**(-1 , -1)**	**(0 , 0)**

Nash Equilibrium

- Instead of hoping to find strategies that players would play *irrespective of what other players play,* we want to find strategies that players would play *given what other players play.*

- Nash Equilibrium
  - A strategy profile $\vec{s}$ is in Nash equilibrium if s_i is the best action for player i given that other players are playing $\vec{s}_{-i}$

$$u_i(s_i, \vec{s}_{-i}) \geq u_i(s_i', \vec{s}_{-i}), \forall s_i'$$

Recap: Prisoner's Dilemma

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

- Nash equilibrium?

- (Dominant strategies)

Recap: Microsoft vs Startup

Microsoft \ Startup	Enter	Stay Out
Enter	(2 , -2)	(4 , 0)
Stay Out	(0 , 4)	(0 , 0)

- Nash equilibrium?
- (Iterated elimination of strongly dominated strategies)

Recap: Attend or Not

Students \ Professor	Attend	Be Absent
Attend	(3 , 1)	(-1 , -3)
Be Absent	(-1 , -1)	(0 , 0)

- Nash equilibria?

- Lack of predictability

Example: Rock-Paper-Scissor

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

- Pure Nash equilibrium?

Nash's Beautiful Result

- Theorem: Every normal form game admits a mixed-strategy Nash equilibrium.

- What about Rock-Paper-Scissor?

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

Indifference Principle

- If the mixed strategy of player i in a Nash equilibrium has support T_i, the expected payoff of player i from each $s_i \in T_i$ must be identical.

- Derivation of rock-paper-scissor on the board.

Stag-Hunt

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

- Game
  - Stag requires both hunters, food is good for 4 days for each hunter.
  - Hare requires a single hunter, food is good for 2 days
  - If they both catch the same hare, they share.

- Two pure Nash equilibria: (Stag,Stag), (Hare,Hare)

Stag-Hunt

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

- Two pure Nash equilibria: (Stag,Stag), (Hare,Hare)
  - Other hunter plays "Stag" → "Stag" is best response
  - Other hunter plays "Hare" → "Hare" is best reponse

- What about mixed Nash equilibria?

Stag-Hunt

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

- Symmetric: $s \rightarrow \{$Stag w.p. p, Hare w.p. $1-p\}$
- Indifference principle:
  - *Given the other hunter plays* s*, equal* $\mathbb{E}$*[reward] for Stag and Hare*
  - $\mathbb{E}[\text{Stag}] = p * 4 + (1-p) * 0$
  - $\mathbb{E}[\text{Hare}] = p * 2 + (1-p) * 1$
  - Equate the two $\Rightarrow p = 1/3$

Extra Fun 1: Cunning Airlines

- Two travelers lose their luggage.
- Airline agrees to refund up to \$100 to each.
- Policy: Both travelers would submit a number between 2 and 99 (inclusive).
  - If both report the same number, each gets this value.
  - If one reports a lower number (s) than the other (t), the former gets s+2, the latter gets s-2.

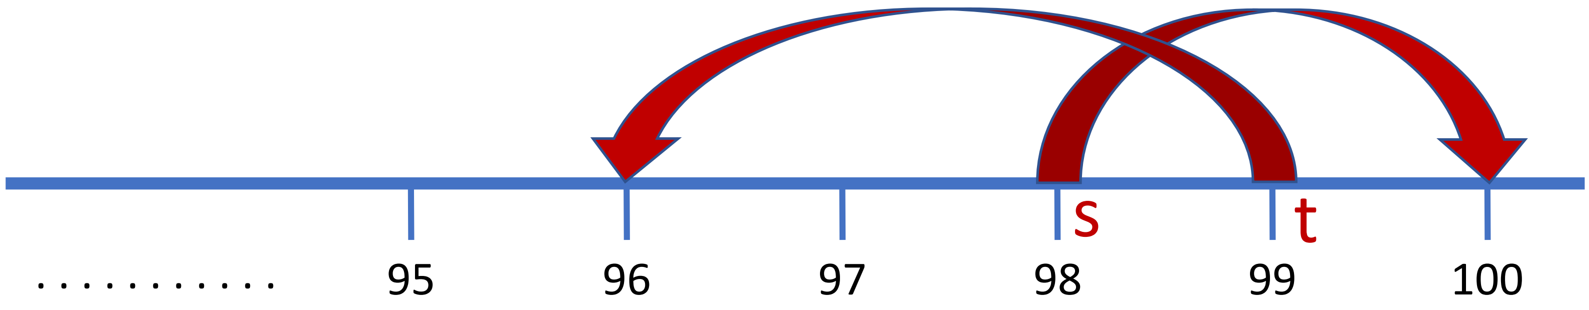

Extra Fun 2: Ice Cream Shop

- Two brothers, each wants to set up an ice cream shop on the beach ([0,1]).
- If the shops are at s, t (with $s \leq t$)
  - The brother at s gets $\left[0, \frac{s+t}{2}\right]$, the other gets $\left[\frac{s+t}{2}, 1\right]$

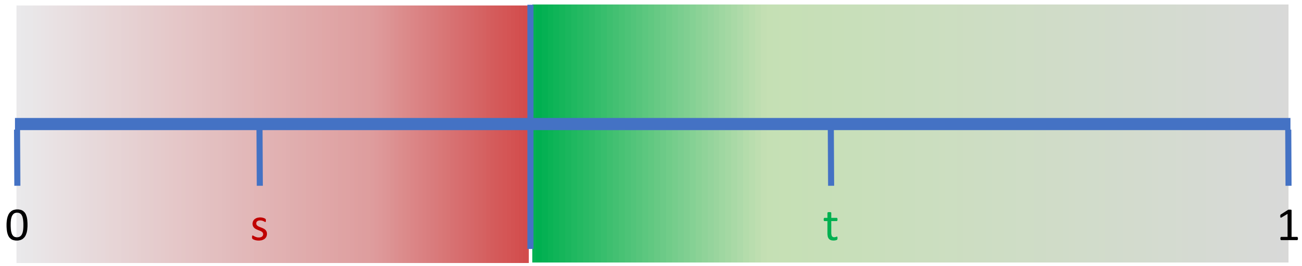

Nash Equilibria: Critique

- Noncooperative game theory provides a framework for analyzing rational behavior.
- But it relies on many assumptions that are often violated in the real world.
- Due to this, human actors are observed to play Nash equilibria in some settings, but play something far different in other settings.

Nash Equilibria: Critique

- Assumptions:
  - Rationality is common knowledge.
    - All players are rational.
    - All players know that all players are rational.
    - All players know that all players know that all players are rational.
    - … [Aumann, 1976]
    - Behavioral economics
  - Rationality is perfect = "infinite wisdom"
    - Computationally bounded agents
  - Full information about what other players are doing.
    - Bayes-Nash equilibria

Nash Equilibria: Critique

- Assumptions:
  - No binding contracts.
    - Cooperative game theory
  - No player can commit first.
    - Stackelberg games (will study this in a few lectures)
  - No external help.
    - Correlated equilibria
  - Humans reason about randomization using expectations.
    - Prospect theory

Nash Equilibria: Critique

- Also, there are often multiple equilibria, and no clear way of “choosing” one over another.

- For many classes of games, finding a single equilibrium is provably hard.
  - Cannot expect humans to find it if your computer cannot.

Nash Equilibria: Critique

- Conclusion:
  - For human agents, take it with a grain of salt.
  - For AI agents playing against AI agents, perfect!

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 $e \in P_i$

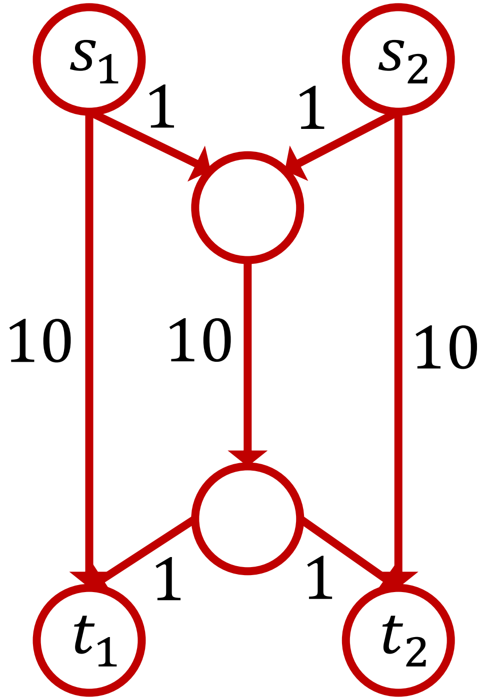

Cost Sharing Game

- Given strategy profile $\vec{P}$, cost $c_i(\vec{P})$ 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 that $C(\vec{P}) = \sum_{e \in E(\vec{P})} c_e$, where $E(\vec{P})$={edges taken in $\vec{P}$ by at least one player}
- 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?

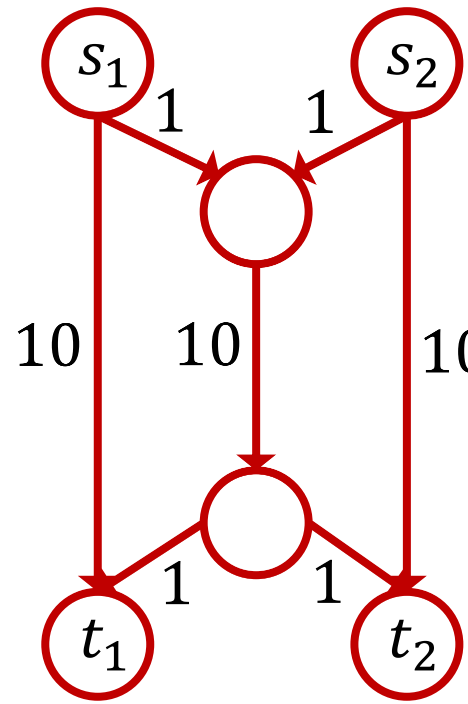

Cost Sharing: Simple Example

- 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

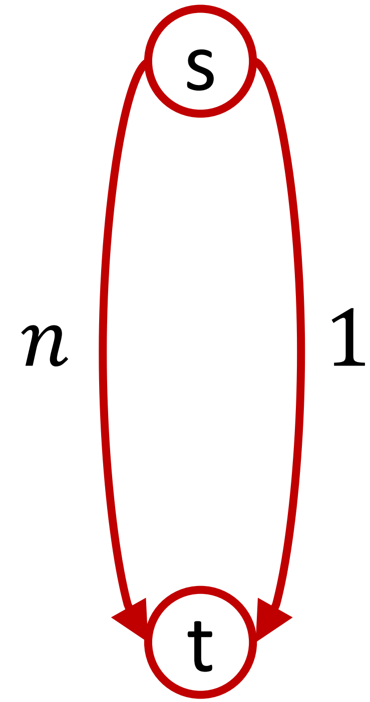

- In this game, price of anarchy $\geq n$
- We can show that for all cost sharing games, price of anarchy $\leq n$

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?

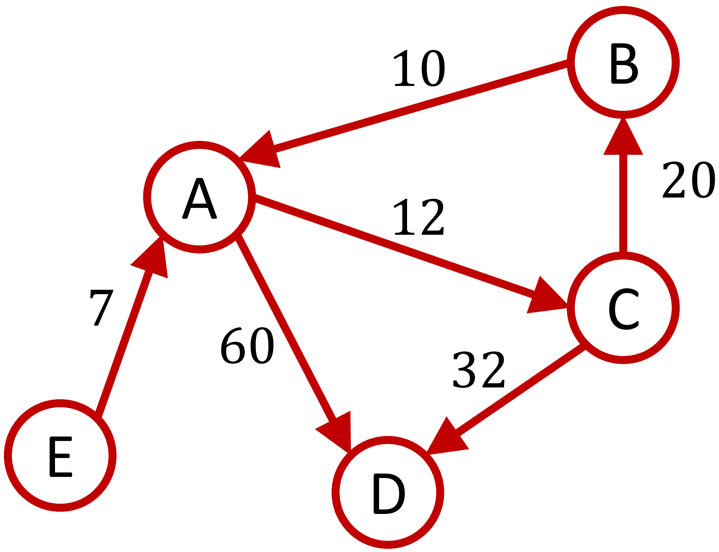

10 players: $E \rightarrow C$
27 players: $B \rightarrow D$
19 players: $C \rightarrow D$

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: $\Phi : \prod_i S_i \rightarrow \mathbb{R}_+$
  - 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 $\vec{P}$

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

- 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.

- 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.

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

Social cost

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

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

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

Potential minimizing eq.

Social optimum

Potential Minimizing Eq.

- Potential minimizing equilibrium gives $O(\log n)$ approximation to the social optimum
  - Price of stability is $O(\log n)$
    - ∃ 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:

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

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

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

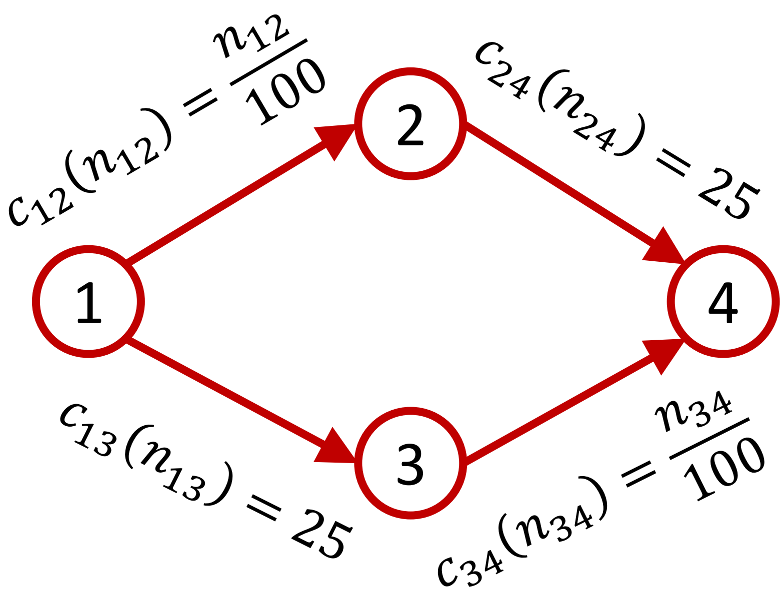

The Braess' Paradox

- 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

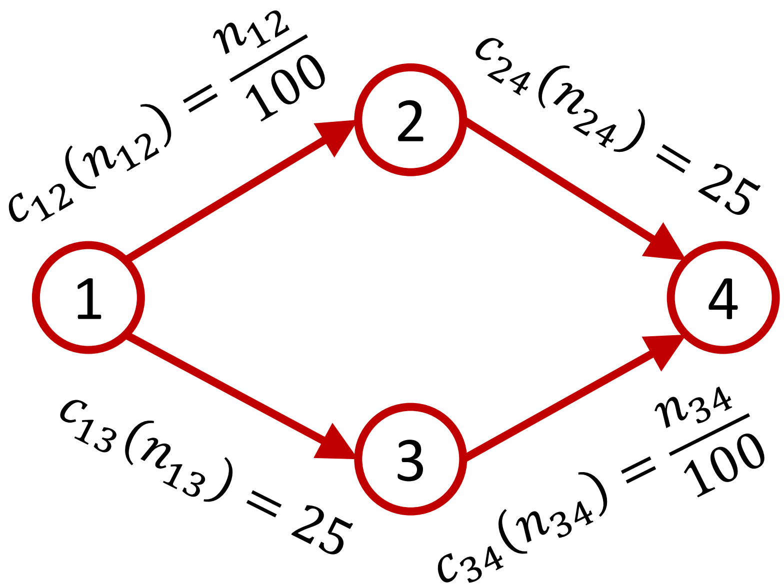

The Braess' Paradox

- 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!

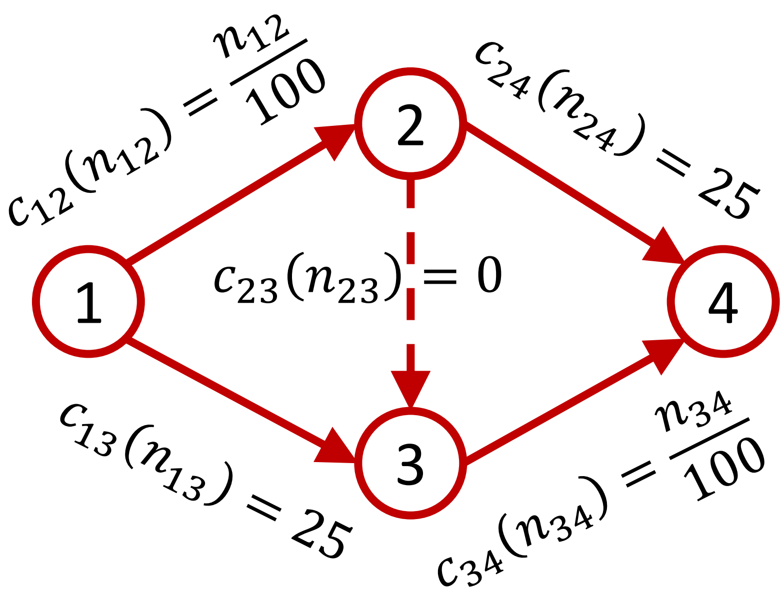

The Braess' Paradox

- 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 $= 40$!

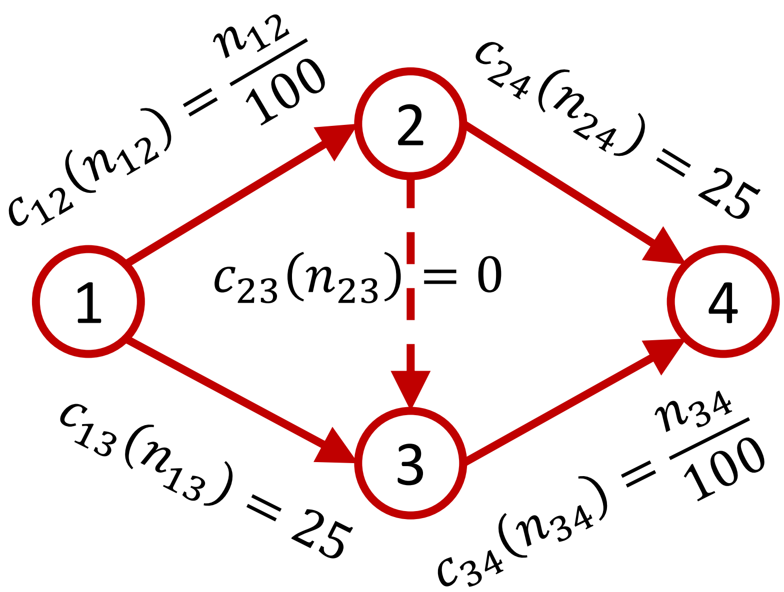

The Braess' Paradox

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