CSC304
Algorithmic Game Theory & Mechanism Design

Nisarg Shah

Recap

- Iterated elimination
  - Even when no player has a dominant action, we can iteratively eliminate dominated actions of players, which can make some previously undominated actions of other players dominated
  - Two versions depending on strict/weak domination

- Nash equilibria
  - Pure Nash equilibria via best response diagram
  - Mixed Nash equilibria via the indifference principle

Nash Equilibrium

- 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 $\vec{s}_{-i}$

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

No quantifier on $\vec{s}_{-i}$

  - Each player's strategy is only best *given* the strategies of others, and not *regardless*.
  - You can't reason about a single player in isolation. You can only say whether you're in a NE after seeing the entire strategy profile.

Recap: Attend or Not

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

- Pure Nash equilibria
  - (Attend, Attend)
  - (Be Absent, Be Absent)

- Not pure Nash equilibria
  - (Attend, Be Absent)
  - (Be Absent, Attend)

Nash's Beautiful Result

- Nash's Theorem:
  - Every normal form game has at least one (possibly mixed) Nash equilibrium.

- The Indifference Principle
  - *If* $(s_i, \vec{s}_{-i})$ *is a Nash equilibrium and* s_i *assigns a positive probability to the set of actions* T_i *of player i, then…*

…for all actions $a_i, a_i' \in T_i$ *and* $a_i'' \notin T_i$

Reward to i for playing a_i w.p. 1 when others play $\vec{s}_{-i}$

$$u_i(a_i, \vec{s}_{-i}) = u_i(a_i', \vec{s}_{-i}) \geq u_i(a_i'', \vec{s}_{-i})$$

Must be indifferent between actions in T_i

Must prefer actions in T_i to any others

Applying Indifference Principle

- Let S_1 and S_2 denote the set of actions of players 1 and 2

- For every $T_1 \subseteq S_1$ ($T_1 \neq \emptyset$) and $T_2 \subseteq S_2$ ($T_2 \neq \emptyset$)
  - Write generic strategies s_1 and s_2 randomizing over T_1 and T_2
  - Apply the indifference principle to player 1 to solve for s_2
  - Apply the indifference principle to player 2 to solve for s_1
    - Sometimes you obtain multiple (or even infinitely many) solutions, in which case all of them are mixed Nash equilibria

Example

Player 2 / Player 1	L	R
T	(5 , 4)	(0 , 2)
M	(2 , 0)	(4 , 1)
B	(3, 20)	(1, 50)

- Case of S_1 = {T,M} and S_2 = {L,R}
  - Let s_1 = (T,M,B) with probabilities $(p, 1-p, 0)$, where $p \in (0,1)$
  - Let s_2 = (L,R) with probabilities $(q, 1-q)$, where $q \in (0,1)$

  - We want to solve for possible values of p and q (if any) by applying the indifference principle

Example

Player 2 / Player 1	L	R
T	(5 , 4)	(0 , 2)
M	(2 , 0)	(4 , 1)
B	(3, 20)	(1, 50)

- Indifference principle to player 1
  - $u_1(T, s_2) = u_1(M, s_2) \geq u_1(B, s_2)$
  - $q \cdot 5 + (1-q) \cdot 0 = q \cdot 2 + (1-q) \cdot 3 \geq q \cdot 3 + (1-q) \cdot 1$
  - $q = {}^4/_7$ works!
    - If player 2 plays (L,R) w.p. $({}^4/_7, {}^3/_7)$, then playing (any) randomization between T and M would be best response for player 1

Example

Player 1 \ Player 2	L	R
T	(5 , 4)	(0 , 2)
M	(2 , 0)	(4 , 1)
B	(3, 20)	(1, 50)

- Indifference principle to player 2
  - $u_2(s_1, L) = u_2(s_1, R)$
  - $p \cdot 4 + (1 - p) \cdot 0 = p \cdot 2 + (1 - p) \cdot 1$
  - $p = {}^1/_3$ works!
    - If player 1 plays (T,M,B) w.p. $({}^1/_3\,, {}^2/_3\,, 0)$, then playing (any) randomization between L and R would be best response for player 2

Exercise: Rock-Paper-Scissor

- Exercise: Solve for the "fully mixed" case where…
  - P1 plays (R,P,S) w.p. $(p, q, 1 - p - q)$, where $p > 0, q > 0, p + q < 1$
  - P2 plays (R,P,S) w.p. $(x, y, 1 - x - y)$, where $x > 0, y > 0, x + y < 1$
  - Apply the indifference principle to P1 to solve for x and y
  - Use symmetry to argue that the same calculations hold for p and q

- Exercise: Check that other cases yield no equilibria

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)

Exercise: Inspect Or Not

Driver \ Inspector	Inspect	Don't Inspect
Pay Fare	(-10 , -1)	(-10 , 0)
Don't Pay Fare	(-90 , 29)	(0 , -30)

- Game:
  - Fare = 10
  - Cost of inspection = 1
  - Fine if fare not paid = 30
  - Total cost to driver if caught = 90

- Pure and mixed Nash equilibria?

Exercise: 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.

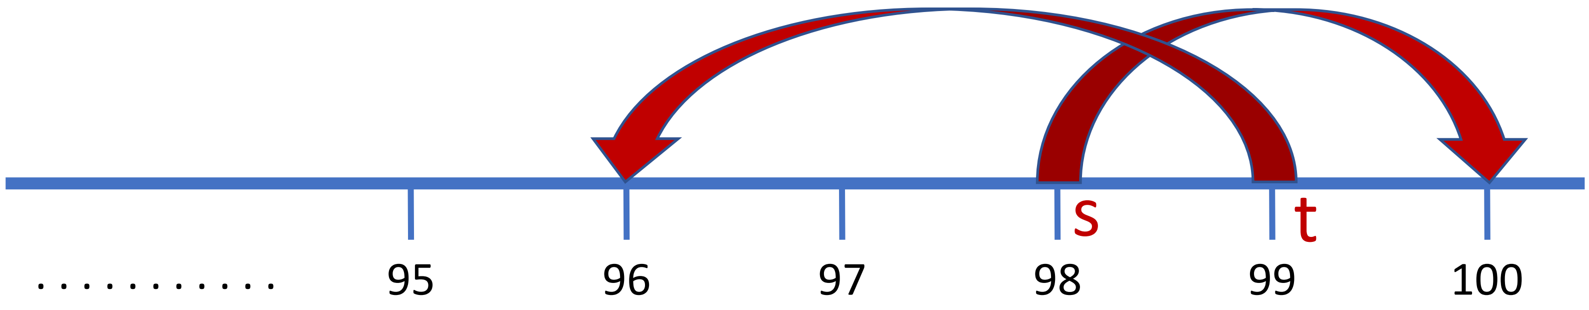

Exercise: Ice Cream Shop

- Two brothers, each wants to set up an ice cream shop on the beach [0,1]
- Reward structure:
  - If the shops are at s, t (with $s \leq t$), the brother at s gets the customers in $\left[0, \frac{s+t}{2}\right]$ and the other one gets the customers in $\left[\frac{s+t}{2}, 1\right]$
  - Reward equals the length of the interval

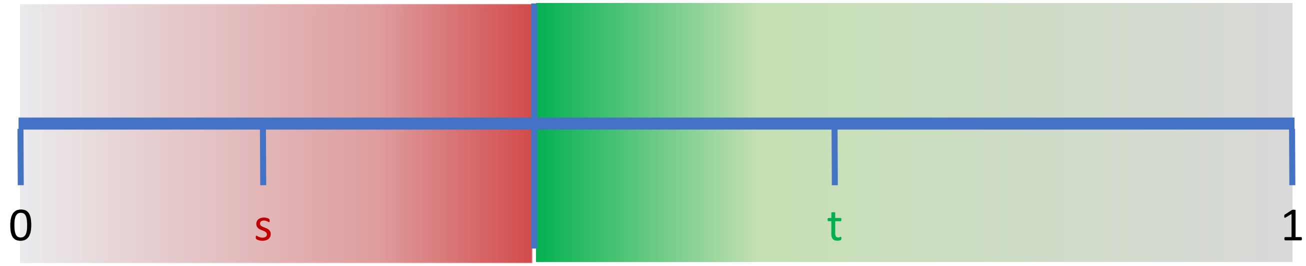

Computational Complexity

- Pure Nash equilibria
  - Existence: Checking the existence of a pure Nash equilibrium can be NP-hard.
  - Computation: Computing a pure NE can be PLS-complete, even in games in which a pure NE is guaranteed to exist.

- Mixed Nash equilibria
  - Existence: Always exist due to Nash's theorem
  - Computation: Computing a mixed NE is PPAD-complete.

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 even a single Nash 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!

Prices of Anarchy & Stability

Worst and Best Nash Equilibria

- What can we say after we identify all Nash equilibria?
  - Compute how "good" they are in the best/worst case

- How do we measure "social good"?
  - Game with only rewards?
    Higher total reward of players = more social good
  - Game with only penalties?
    Lower total penalty to players = more social good
  - Game with rewards and penalties?
    No clear consensus…

Price of Anarchy and Stability

- 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

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

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

Revisiting Stag-Hunt

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

- Max total reward $= 4 + 4 = 8$
- Three equilibria
  - (Stag, Stag) : Total reward = 8
  - (Hare, Hare) : Total reward = 2
  - ($^1/_3$ Stag – $^2/_3$ Hare, $^1/_3$ Stag – $^2/_3$ Hare)
    - Total reward = $\frac{1}{3} * \frac{1}{3} * 8 + \left(1 - \frac{1}{3} * \frac{1}{3}\right) * 2 \in (2,8)$
- Price of stability? Price of anarchy?

Revisiting Prisoner's Dilemma

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

- Min total cost $= 1 + 1 = 2$
- Only equilibrium:
  - (Betray, Betray) : Total cost $= 2 + 2 = 4$

- 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 $s_i \to t_i$ 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 $e \in P_i$

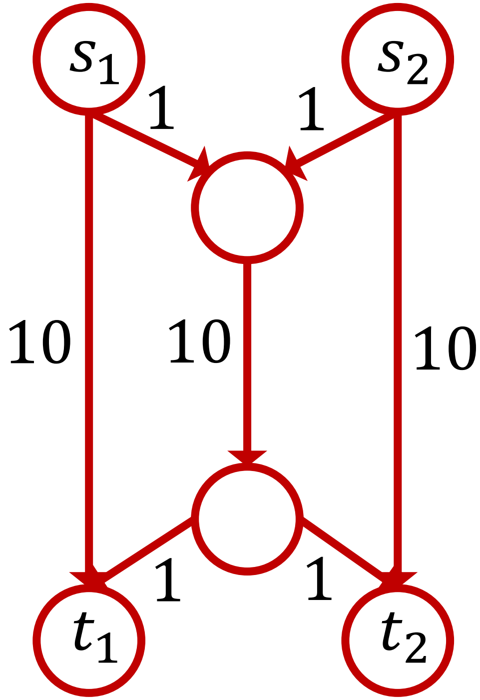

Cost Sharing Game

- 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?

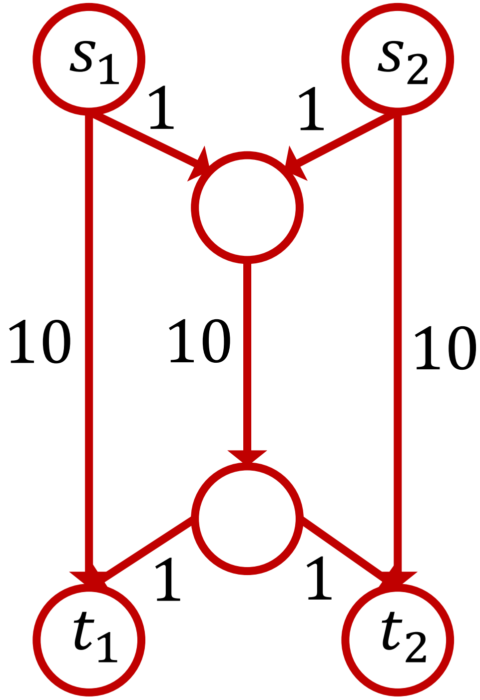

Cost Sharing Game

- In the example on the right:
  - What if both players take direct paths?
  - What if both take middle paths?
  - What if one player takes direct path and the other takes middle path?

- Pure Nash equilibria?

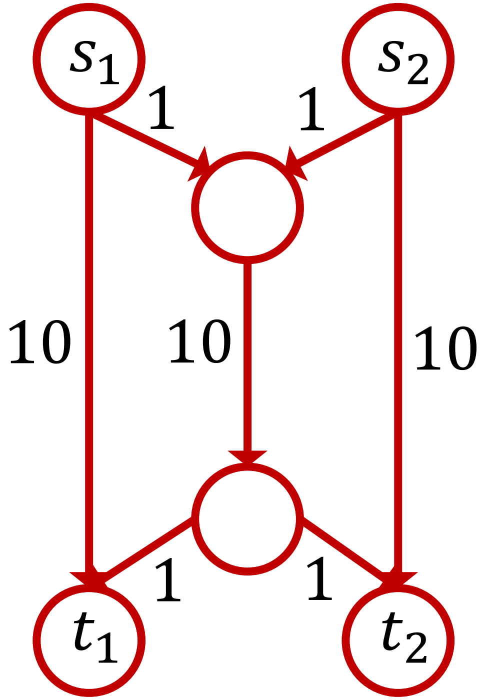

Cost Sharing: Simple Example

- 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
- Price of stability: 1
- Price of anarchy: n
  - We can show that price of anarchy $\leq n$ in *every* cost-sharing game!

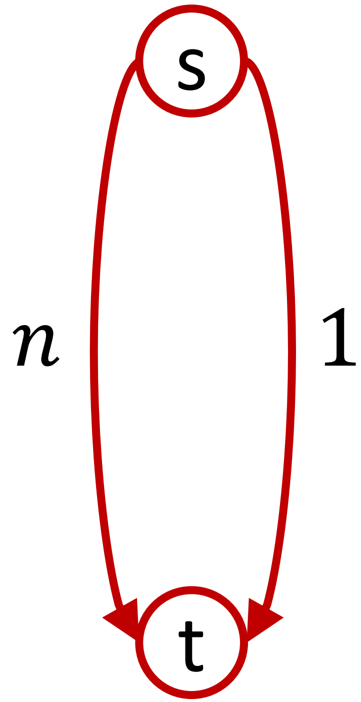

Cost Sharing: PoA

- 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
  - Every cost-sharing game: PoA $\leq n$
  - Example game with PoA $= n$
  - Bound of n is tight.

- Price of stability?
  - In the previous game, it was 1.
  - In general, it can be higher. How high?
  - We'll answer this after a short detour.

Cost Sharing

- Nash's theorem shows existence of a mixed NE.
  - Pure NE may not always exist in general.
- But in both cost-sharing games we saw, there was a PNE.
  - What about a more complex game like the one on the right?

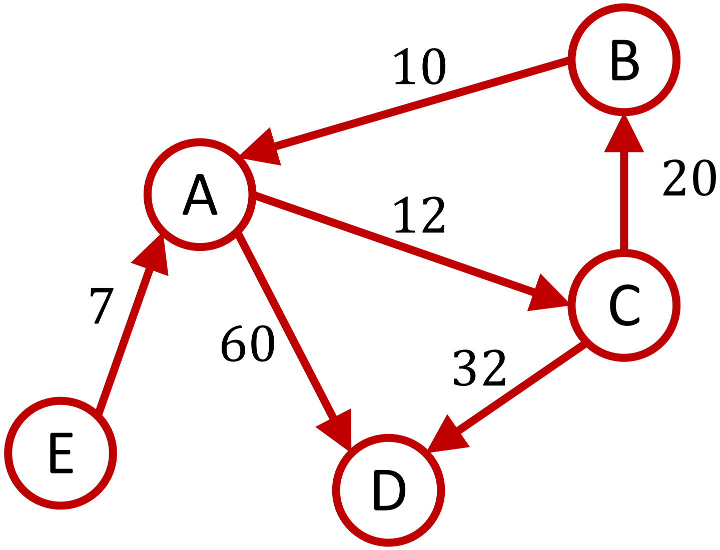

10 players: $E \to C$
27 players: $B \to D$
19 players: $C \to D$