CSC304
Algorithmic Game Theory & Mechanism Design

Nisarg Shah

Introduction

- **Instructor:** Nisarg Shah (~nisarg, SF 2301C)
- **Course email:** csc304-2022-09@cs.toronto.edu
  - Please use this rather than my personal email as we may have a co-instructor added to the course.
- **TAs:** Soroush Ebadian, Mohamad Latifian, Devansh Shringi
- **Lectures:** Tue, 3-5pm, BA 1190
- **Tutorials:** Thu, 4-5pm, division by birth month
  Jan-Apr → BA 2135
  May-Aug → BA 2135
  Sep-Dec → BA 2135
- **Office hours:** Poll!

There was no tutorial on Sep 8

First tutorial will be this Thursday, Sep 15.

Course Information

- Course Page:
  `www.cs.toronto.edu/~nisarg/teaching/304f22/`

- Discussion Board:
  `piazza.com/utoronto.ca/fall2022/csc304`

- Assignments & Grades:
  `https://markus.teach.cs.toronto.edu/2022-09`

  - LaTeX preferred, scans are OK!

Course Organization

- Three (roughly equal) parts:
  - Game theory
  - Mechanism design with money
  - Mechanism design without money

- A homework and a midterm for each part

- Final exam = “third midterm” + a section on entire syllabus

Textbook

- Not really.
  - Slides will be your main reference.

- But…but…I want a textbook?
  - OK…
  - Book by Prof. David Parkes at Harvard
    - In preparation…
    - Closely follows the syllabus structure
    - Available from my webpage (username/password on Piazza)
  - Other good books mentioned in the handout

Grading Policy

- 3 homeworks * 15% = 45%
- 3 midterms * 15% = 45%
- Final exam (entire syllabus) = 10%
  - Final exam: third midterm + entire syllabus = 15+10 = 25%

Other Policies

- Collaboration
  - Assignments in groups of size up to 3
  - You can form your own groups on MarkUs
  - Free to discuss with other groups or read online material.
  - Must write solutions in your own words (easier if you do not take any pictures/notes from the discussions)
- Citation
  - For each question, must cite the peer (write the name) or the online sources (provide links) referred, if any.
  - Failing to do this is plagiarism!

Other Policies

- "No Garbage" Policy
  - Borrowed from: Prof. Allan Borodin (citation!)
  1. Partial marks for viable approaches
  2. Zero marks if the answer makes no sense
  3. 20% marks if you admit to not knowing how to solve (10% if you do not write any admission and just leave it blank)
     - 20% > 0% !!
- Applies to assignments+exams
  - To questions and even to subquestions
  - Doesn't apply to bonus questions

Other Policies

- Late Days
  - 3 late days total across 3 homeworks
  - At most 2 late days for a single homework
  - Covers legitimate reasons such as illness, University activities, etc.

Enough with the boring stuff.

What will we study?

Why will we study it?

What is this course about?

- Game Theory and Mechanism Design
  - Topics from microeconomics
- + Computer Science:
  - Algorithmic Game Theory (AGT)
  - Algorithmic Mechanism Design (AMD)

Game Theory

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

- We focus on noncooperative games
  - No external force or agencies forming coalitions

Example: Prisoner's Dilemma

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

- What Sam thinks:
  - If John is going to stay silent…
    - Better for me to betray (my reward: 0)
    - Than for me to stay silent (my reward: -1)
  - If John is going to betray…
    - Better for me to betray (my reward: -2)
    - Than for me to stay silent (my reward: -3)

Makes sense only to betray

John thinks the same way

That’s cute.

But is this really useful in the real world?

Security Games

Deploying “patrol units” to protect infrastructure targets, prevent smuggling, save wildlife…

LA Metro

LAX

Staten Island Ferry

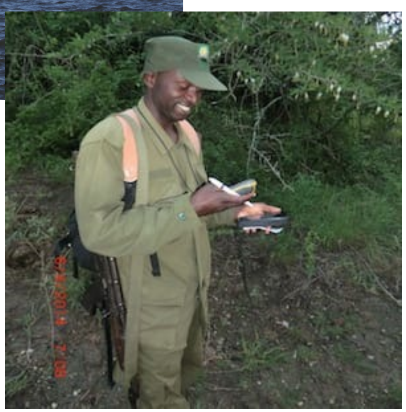

Ugandan Forest

Image Courtesy: Teamcore

Security Games

- n targets
- Player 1: Attacker
  - Actions: attack a target
- Player 2: Defender
  - Actions: protect k ($< n$) targets at a time
  - $\binom{n}{k}$ actions – exponential!
- Attacker can observe ⇒ need to randomize
- Large games ⇒ need fast algorithms

Mechanism Design

- Design the rules of the game
- A principal in the system
  - Wants the n rational agents to behave "nicely"
- Decides the rewards (or penalties) as a function of actions to incentivize the desired behavior
  - Often the desired behavior is unclear
  - E.g., want agents to reveal their true preferences

Mechanism Design

- With money
  - Principal can "charge" the agents (require payments)
  - Helps significantly
  - Example: auctions

- Without money
  - Monetary transfers are not allowed
  - Incentives must be balanced otherwise
  - Often impossible without sacrificing the objective a little
  - Example: elections, kidney exchange

Example: Auction

Objective: The one who really needs it more should have it.

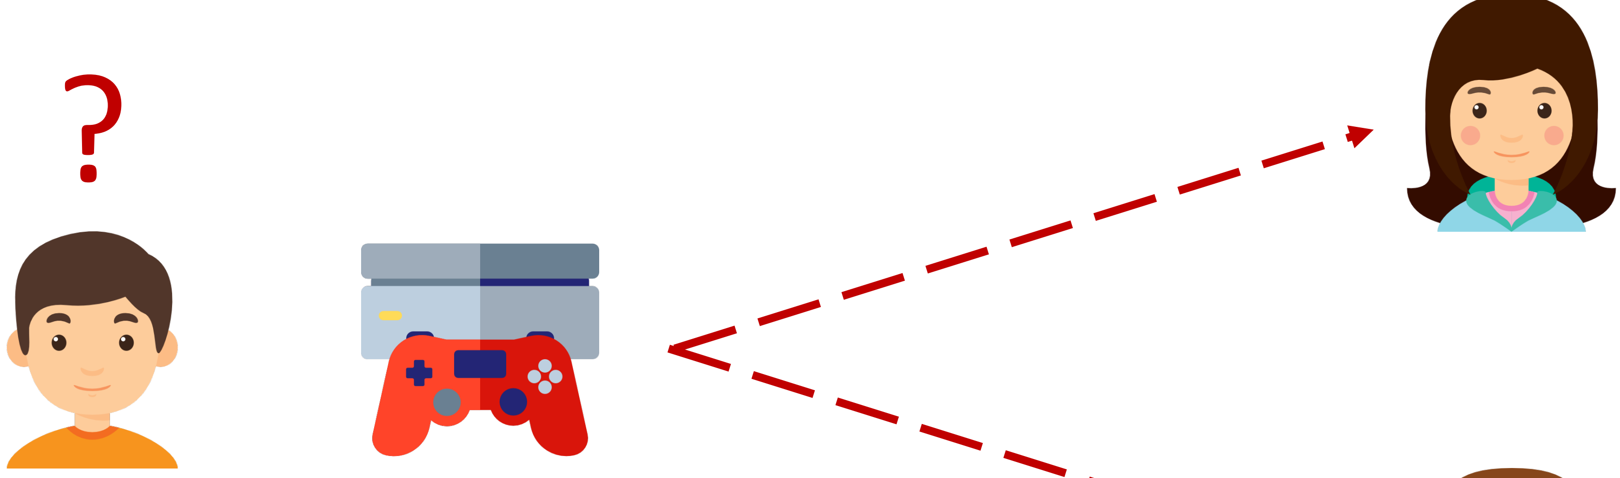

Rule 1: Each would tell me his/her value.
I'll give it to the one with the higher value.

Image Courtesy: Freepik

Example: Auction

Objective: The one who really needs it more should have it.

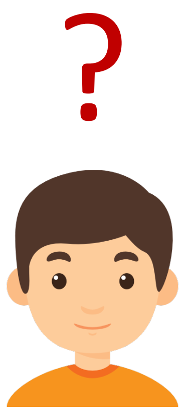

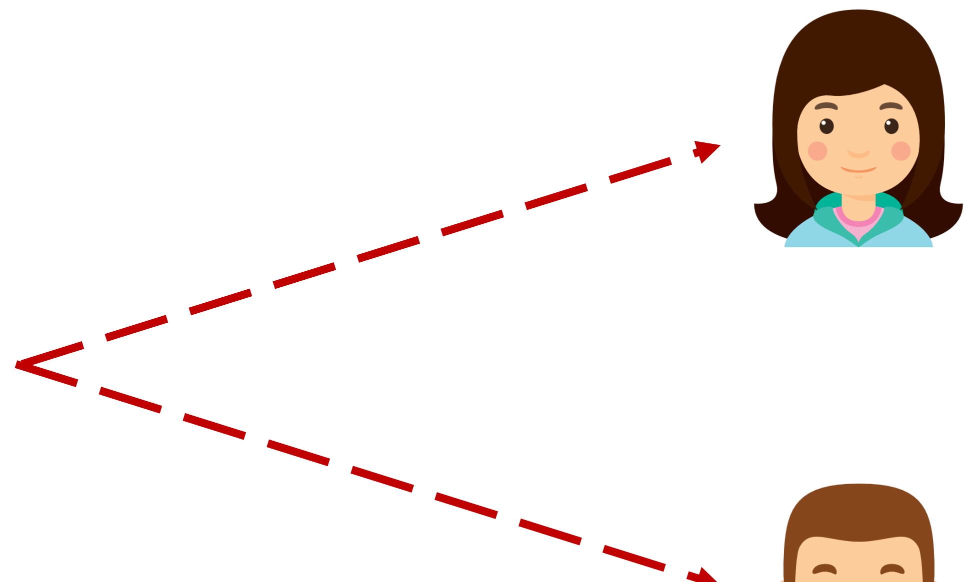

Rule 2: Each would tell me his/her value. I'll give it to the one with the higher value, but they have to pay me that value.

Image Courtesy: Freepik

Example: Auction

Objective: The one who really needs it more should have it.

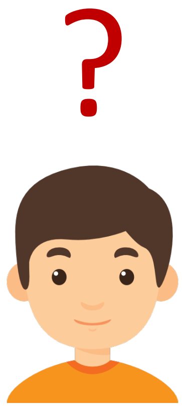

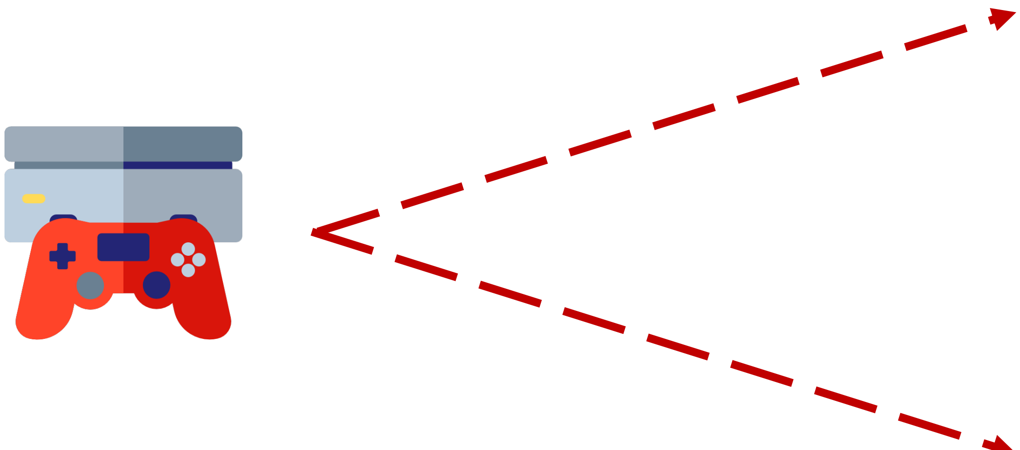

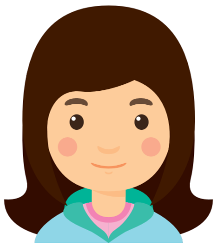

Can I make it easier so that each can just truthfully tell me how much they value it?

Image Courtesy: Freepik

Real-World Applications

- Auctions form a significant part of mechanism design with money
- Auctions are ubiquitous in the real world!
  - A significant source of revenue for many large organizations (including Facebook and Google)
  - Often run billions of tiny auctions everyday
  - Need the algorithms to be fast

Example: Facility Location

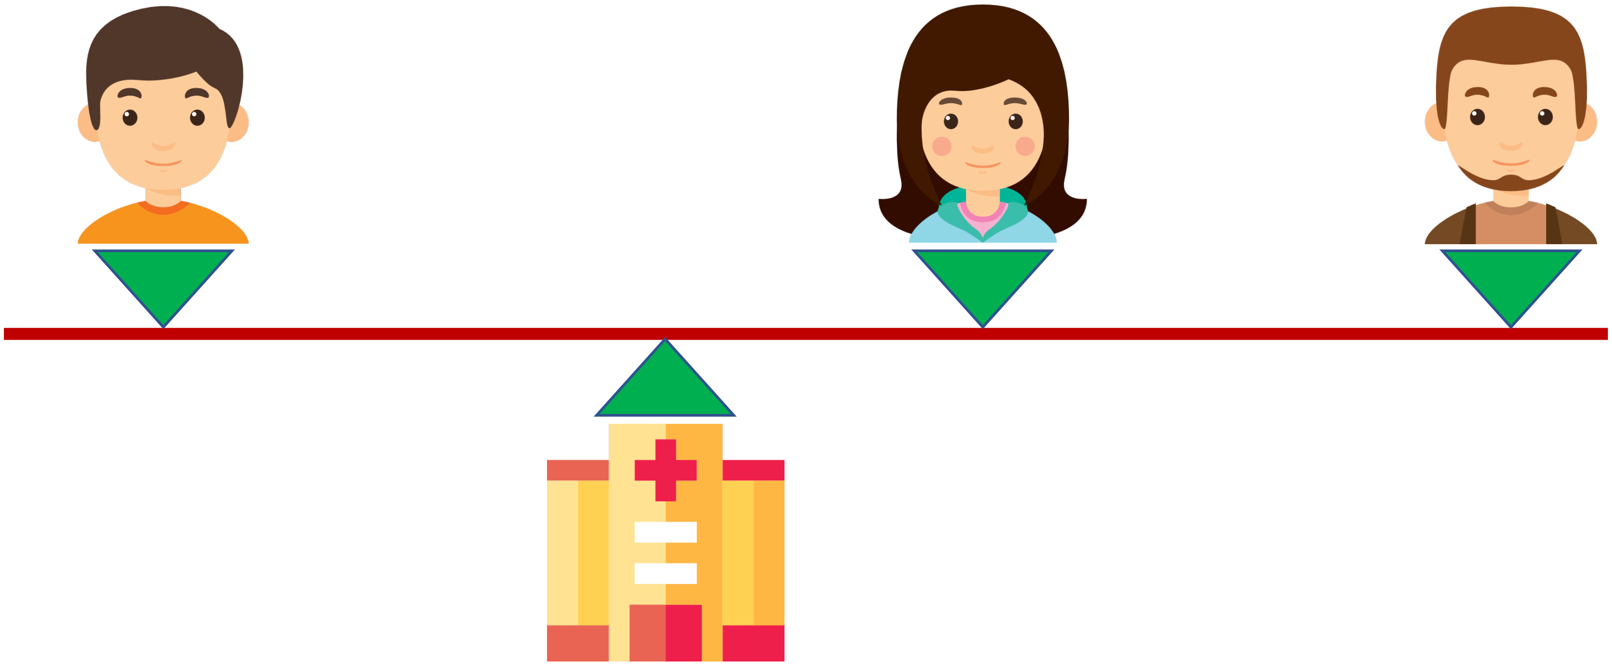

Cost to each agent: Distance from the hospital

Objective: Minimize the sum of costs

Constraint: No money

Image Courtesy: Freepik

Example: Facility Location

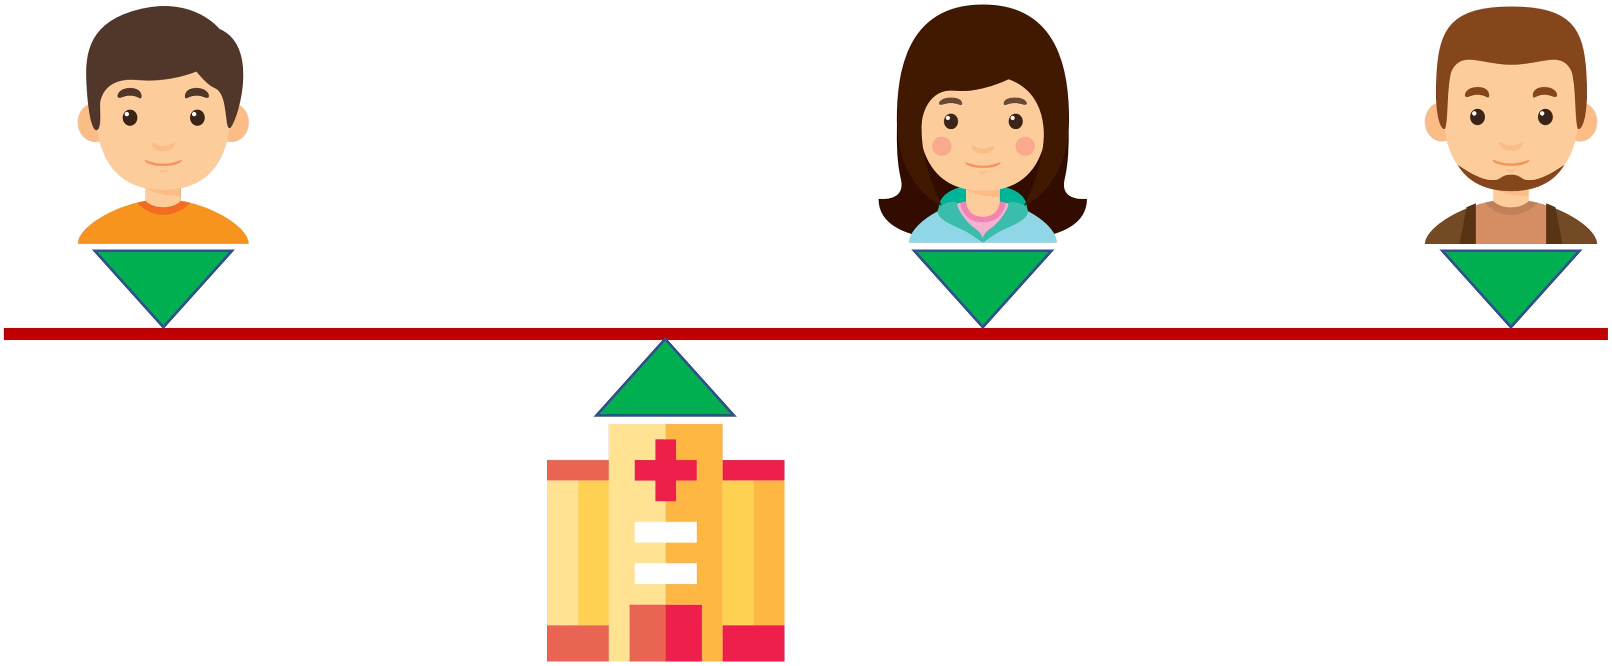

Q: What is the optimal hospital location?

Q: If we decide to choose the optimal location, will the agents really tell us where they live?

Image Courtesy: Freepik

Example: Facility Location

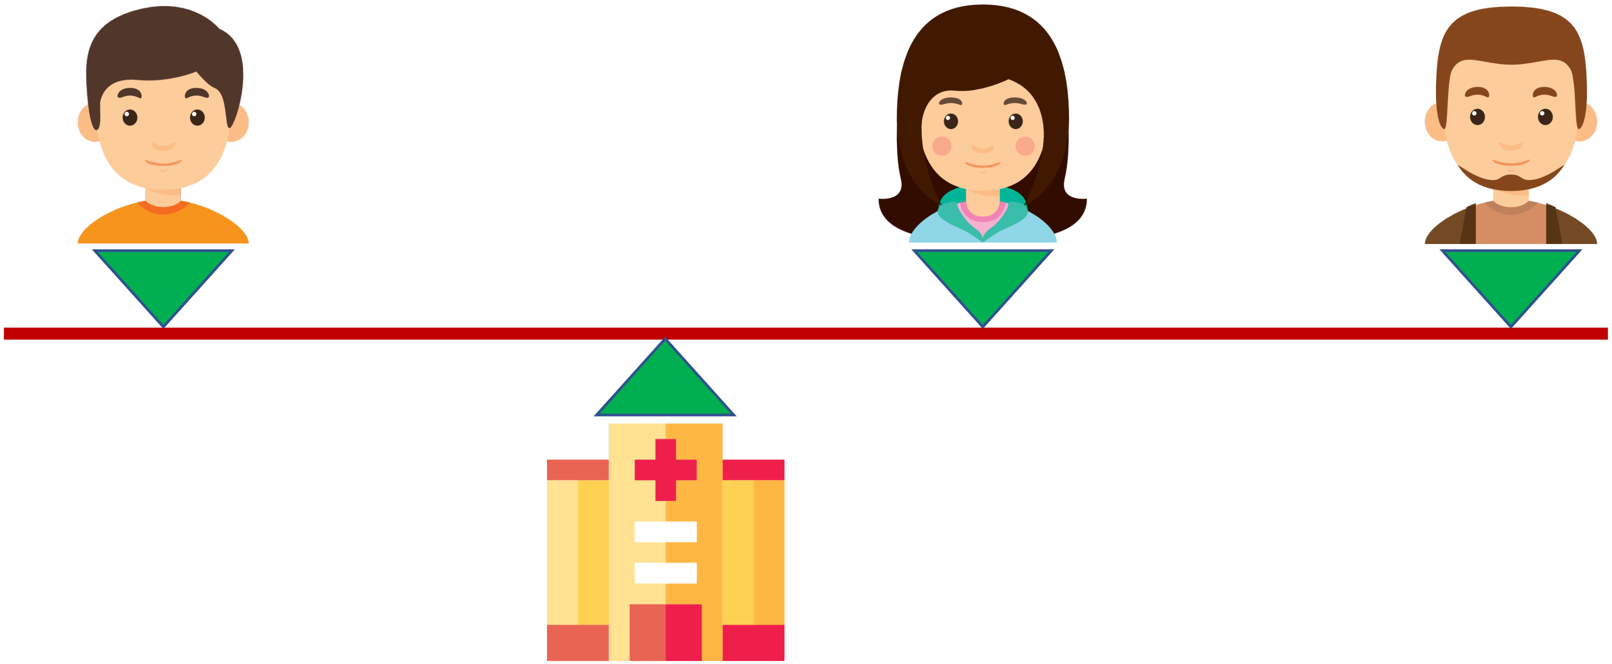

Cost to each agent: Distance from the hospital

Objective: Minimize the maximum cost

Constraint: No money

Image Courtesy: Freepik

Example: Facility Location

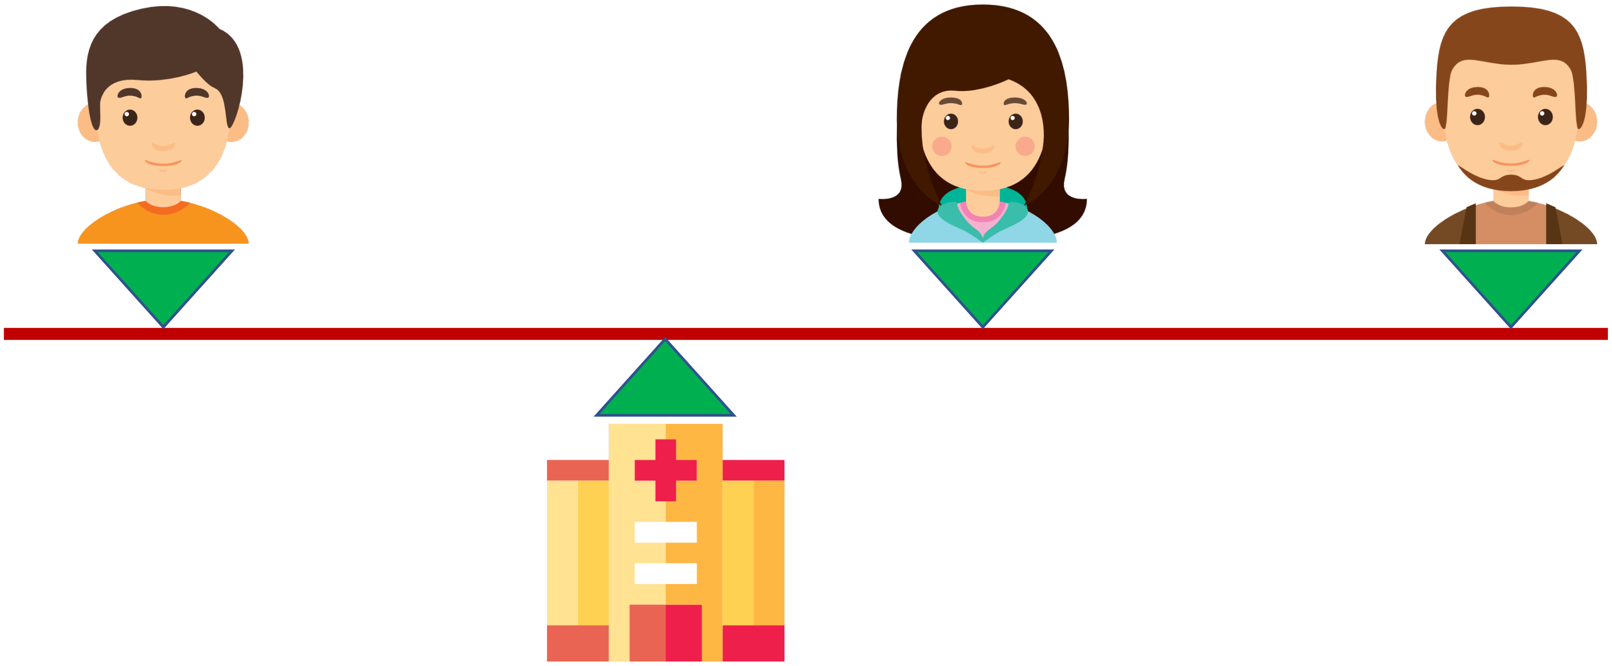

Q: What is the optimal hospital location?

Q: If we decide to choose the optimal location, will the agents really tell us where they live?

Image Courtesy: Freepik

Mechanism Design w/o Money

- Truth-telling is not the only possible desideratum
  - Fairness
  - Stability
  - Efficiency
  - …

- Consequently, many subfields of study
  - Fair allocation of resources
  - Stable matching
  - Voting

Real-World Applications


Roth

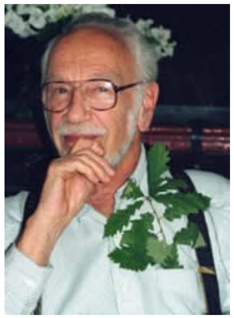
Gale


Shapley

National Resident Matching Program (NRMP)

School Choice (New York, Boston)

Fair Division

spliddit

Voting

ROBOVOTE

Game Theory

Game Theory

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

- We focus on noncooperative games
  - No external force or agencies enforcing coalitions

Normal Form Games

- A set of players $\mathrm{N} = \{1, \dots, n\}$
- A set of actions S
  - Action of player $i \rightarrow s_i$
  - Action profile $\vec{s} = (s_1, \dots, s_n)$

- For each player i, utility function $u_i : S^n \rightarrow \mathbb{R}$
  - Given action profile $\vec{s} = (s_1, \dots, s_n)$, each player i gets reward $u_i(s_1, \dots, s_n)$

Normal Form Games

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

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

$u_{Sam}(Betray, Silent)$

$u_{John}(Betray, Silent)$

s_{Sam} s_{John}

Player Strategies

- Pure strategy
  - Choose an action to play
  - E.g., "Betray"
  - For our purposes, simply an action.
    - In repeated or multi-move games (like Chess), need to choose an action to play at every step of the game based on history.
- Mixed strategy
  - Choose a probability distribution over actions
    - Randomize over pure strategies
    - Every pure strategy is also a mixed strategy
  - E.g., "Betray with probability 0.3, and stay silent with probability 0.7"

Pure strategy domination

- Pure strategy s_i dominates pure strategy s_i' if player i is always "better off" playing s_i than s_i', regardless of the strategies of other players.

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

Example

P2 / P1	b_1	b_2
a_1	(2 , 3)	(4 , 1)
a_2	(2 , 5)	(6 , 3)
a_3	(3 , 1)	(5 , 2)

- P1
  - a_1 vs a_2 ?
  - a_1 vs a_3 ?
  - a_2 vs a_3 ?
- P2
  - b_1 vs b_2 ?

Dominant Pure Strategies

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

- Strict dominance is a strong concept
  - A player who has a strictly dominant strategy has no reason *not* to play it
  - If every player has a strictly dominant strategy, such strategies will very likely dictate the outcome of the game

Example

P1 \ P2	b_1	b_2
a_1	**(2 , 3)**	**(4 , 1)**
a_2	**(2 , 5)**	**(6 , 3)**
a_3	**(3 , 1)**	**(5 , 2)**

- Does either player have a dominant strategy?

Example

P1 \ P2	b_1	b_2	b_3
a_1	**(2 , 3)**	**(4 , 1)**	**(2 , 3)**
a_2	**(2 , 5)**	**(6 , 3)**	**(3 , 5)**
a_3	**(3 , 1)**	**(5 , 2)**	**(4 , 3)**

- How about now?

Example

P1 \ P2	b_1	b_2	b_3
a_1	(2 , 3)	(4 , 1)	(2 , 4)
a_2	(2 , 5)	(6 , 3)	(3 , 6)
a_3	(3 , 1)	(5 , 2)	(4 , 3)

- How about now?

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)**

- Betraying is a 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
- Two variants depending on what we eliminate:
  - Only strictly dominated? Or also weakly dominated?

Iterated Elimination

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

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

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

Iterated Elimination

- More serious: "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!

- In-class poll!

- Recall: We have a unique optimal strategy only if everyone is rational, and everyone thinks everyone is rational, and so on.

Nash Equilibrium

- What if we don't find a unique outcome after iterated elimination of dominated strategies?

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

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: Prisoner's Dilemma

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

- Pure strategy Nash equilibria?

- Food for thought:
  - What is the relation between iterated elimination of weakly/strictly dominated strategies and Nash equilibria?

Recap: Microsoft vs Startup

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

- Pure strategy Nash equilibria?

Recap: Attend or Not

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

- Pure strategy Nash equilibria?