CSC304
Algorithmic Game Theory & Mechanism Design

Nisarg Shah

Announcements

- Assignment 1
  - Due 11:59pm on Saturday, Oct 15
  - You can use up to 2 late days
  - Submit a single PDF named "hwk1.pdf" on MarkUs

- Midterm 1
  - Thursday, Oct 20, 4:10pm – 5:00pm (tutorial slot)
  - In-person
  - EX 100 (Exam Centre)
  - Aid: One 8.5" x 11" sheet of handwritten notes on one side
  - Syllabus: Game theory (first lecture to end of game theory portion in today's lecture)

Stackelberg Games

Recap

- 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

Recap Example

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

- Three outcomes
  - Nash equilibrium: (Up, Left), reward of P1 = 1
  - P1 commits to Down: P2 responds with Right, reward of P1 = 2
  - P1 commits to (0.5 x Up + 0.5 x Down): P2 responds still with Right, reward of P1 = 0.5 x 2 + 0.5 x 3 = 2.5

Stackelberg vs Nash

- Committing first is always better than playing a simultaneous-move game?

- Yes!
  - If (x_1^*, x_2^*) is a NE, P1 is always free to commit to x_1^*, which ensures that P2 will play x_2^* and P1 will get the NE reward
  - P1 may be able to commit to a better strategy than x_1^*

- Applications to security
  - Law enforcement is better off committing to a mixed patrolling strategy and announcing the strategy publicly!

Stackelberg in Zero-Sum

- Recall the minimax theorem:

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

- P1 goes first:
  - P1 chooses maximin strategy x_1^* maximizing $\min_{x_2} (x_1^*)^T A\, x_2$
  - P2 responds with $\operatorname{argmin}_{x_2} (x_1^*)^T A\, x_2$
- P2 goes first:
  - P2 chooses minimax strategy x_1^* minimizing $\max_{x_1} x_1^T A\, x_2^*$
  - P1 responds with $\operatorname{argmax}_{x_1} x_1^T A\, x_2^*$

Minimax Theorem

Both scenarios are identical and equivalent to Nash equilibria.

Stackelberg in General-Sum

- 2-player non-zero-sum game with reward matrices A for P1 and $B \neq -A$ for P2

- What will P1 commit to?

$$\max_{x_1} x_1^T A \; f(x_1)$$

$$\text{where } f(x_1) = \underset{x_2}{\operatorname{argmax}} \; x_1^T B \; x_2$$

- How do we compute this?

Example

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

- Let us separately maximize the reward of P1 in 2 cases:
  - Strategies that cause P2 to play Left
  - Strategies that cause P2 to play Right

- Suppose P1 commits to Up w.p. p, Down w.p. $1 - p$

Example

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

- Strategies that cause P2 to play Left

$$\text{Max} \quad p \cdot 1 + (1-p) \cdot 0$$

such that

$$p \cdot 1 + (1-p) \cdot 0 \geq p \cdot 0 + (1-p) \cdot 1$$

$$p \in [0,1]$$

Reward of P1 assuming P2 plays Left

Causing P2 to play Left

Example

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

- Strategies that cause P2 to play Left

$$\text{Max} \quad p$$
$$\text{such that}$$
$$p \geq 1 - p$$
$$p \in [0,1]$$

Best reward across all strategies where P2 responds with Left = 1

Example

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

- Strategies that cause P2 to play Right

$$\text{Max } \; p \cdot 3 + (1-p) \cdot 2$$

such that

$$p \cdot 1 + (1-p) \cdot 0 \leq p \cdot 0 + (1-p) \cdot 1$$

$$p \in [0,1]$$

Reward of P1 assuming P2 plays Right

Causing P2 to play Right

Example

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

- Strategies that cause P2 to play Right

$$\begin{aligned}&\text{Max} \quad p+2\\&\text{such that}\\&2p \leq 1\\&p \in [0,1]\end{aligned}$$

Best reward across all strategies where P2 responds with Right = 2.5

Example

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

- Since P1 can commit to any strategy…
  - P1 can choose the best among both types of strategies: those that cause P2 to choose Left and those that cause P2 to choose Right
  - Hence, the best possible reward for P1 is the maximum of the two answers

Stackelberg via LPs

NOT IN SYLLABUS

- General algorithm:
  - For each action s_2^* of P2, write a *linear program*
    - Variables: probabilities of P1 playing different actions under a mixed strategy x_1
    - Objective: maximize the reward of P1 when P1 plays x_1 and P2 responds with s_2^*
    - Constraint: s_2^* must be the best response for P2 when P1 plays x_1
  - # linear programs = # actions of P2
    - P1's reward in Stackelberg equilibrium = best answer across all the linear programs
  - Running time: polynomial in the number of actions of P1 and P2

Real-World Applications

NOT IN SYLLABUS

- Security Games
  - Defender (leader) and attacker (follower)
  - Defender assigns patrol units to protect sets of targets, attacker chooses a target to attack
  - Both have different utilities for protecting/attacking different targets
  - Running time polynomial in #actions
    - But #actions exponentially many

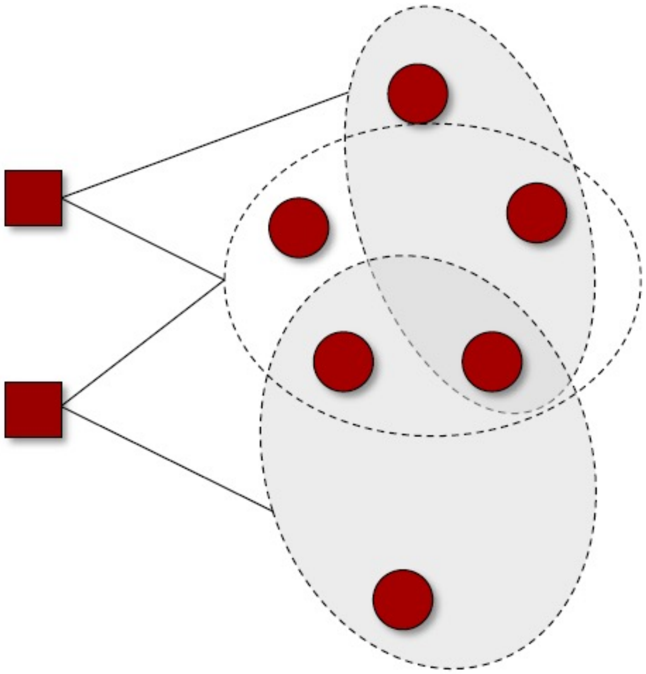

NOT IN SYLLABUS

LAX

Newsweek National News

Subscribe Now | Make Newsweek Your Homepage | Newsletters | RSS

The Element of Surprise

To help combat the terrorism threat, officials at Los Angeles Inter Airport are introducing a bold new idea into their arsenal: randon of security checkpoints. Can game theory help keep us safe?

WEB EXCLUSIVE

By Andrew Murr
Newsweek
Updated: 1:00 p.m. PT Sept 28, 2007

Security forces work the sidewalk

Sept. 28, 2007 - Security officials at Los Angeles International Airport now have a new weapon in their fight against terrorism: complete, baffling randomness. Anxious to thwart future terror attacks in the early stages while plotters are casing the airport, LAX security patrols have begun using a new software program called ARMOR, NEWSWEEK has learned, to make the placement of security checkpoints completely unpredictable. Now all airport security officials have to do is press a button labeled "Randomize," and they can throw a sort of digital cloak of invisibility over where they place the cops' antiterror checkpoints on any given day.

Real-World Applications

NOT IN SYLLABUS

- Protecting entry points to LAX
- Scheduling air marshals on flights
  - Must return home
- Protecting the Staten Island Ferry
  - Continuous-time strategies
- Fare evasion in LA metro
  - Bathroom breaks !!!
- Wildlife protection in Ugandan forests
  - Poachers are not fully rational
- Cyber security

End of Game Theory

Start of Mechanism Design with Money

Mechanism Design with Money

- Design the game structure in order to induce the desired behavior from the agents
- Desired behavior?
  - We will mostly focus on incentivizing agents to truthfully reveal their "private" information
    - Something only the agents know, such as how much value they place on some items
- With money
  - Can pay agents or ask agents for money depending on what the agents report

Mathematical Setup

- A = finite set of outcomes
- Each agent i has a private valuation $v_i : A \to \mathbb{R}$
  - Agent i might report $\tilde{v}_i$ instead of the true v_i
- Mechanism consist of a pair of rules (f, p)
  - Input: reported valuations $\tilde{v} = (\tilde{v}_1, \dots, \tilde{v}_n)$
  - $f(\tilde{v}) \in A$ is the outcome implemented
  - $p(\tilde{v}) = (p_1, \dots, p_n)$ are the payments
    - $p_i(\tilde{v})$ is the amount agent i needs to pay
    - Each agent's payment depends on everyone's reports
- Utility to agent i : $u_i(\tilde{v}) = v_i\big(f(\tilde{v})\big) - p_i(\tilde{v})$

Value minus payment

Mathematical Setup

- Our goal is to design the mechanism (f, p)
  - f is called the social choice function
  - p is called the payment scheme

- Example
  - Suppose we want to sell one item to one of n agents
  - A = set of n outcomes
    - Each corresponds to giving the item to a different agent
  - Agent i values the item at v_i, but may report $\widetilde{v_i}$
    - v_i is the value for receiving the item, value for all other outcomes is 0
  - f takes $\tilde{v}$ as input and decides who gets the item
  - p takes $\tilde{v}$ as input and decides who pays how much

Single-Item Auction

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

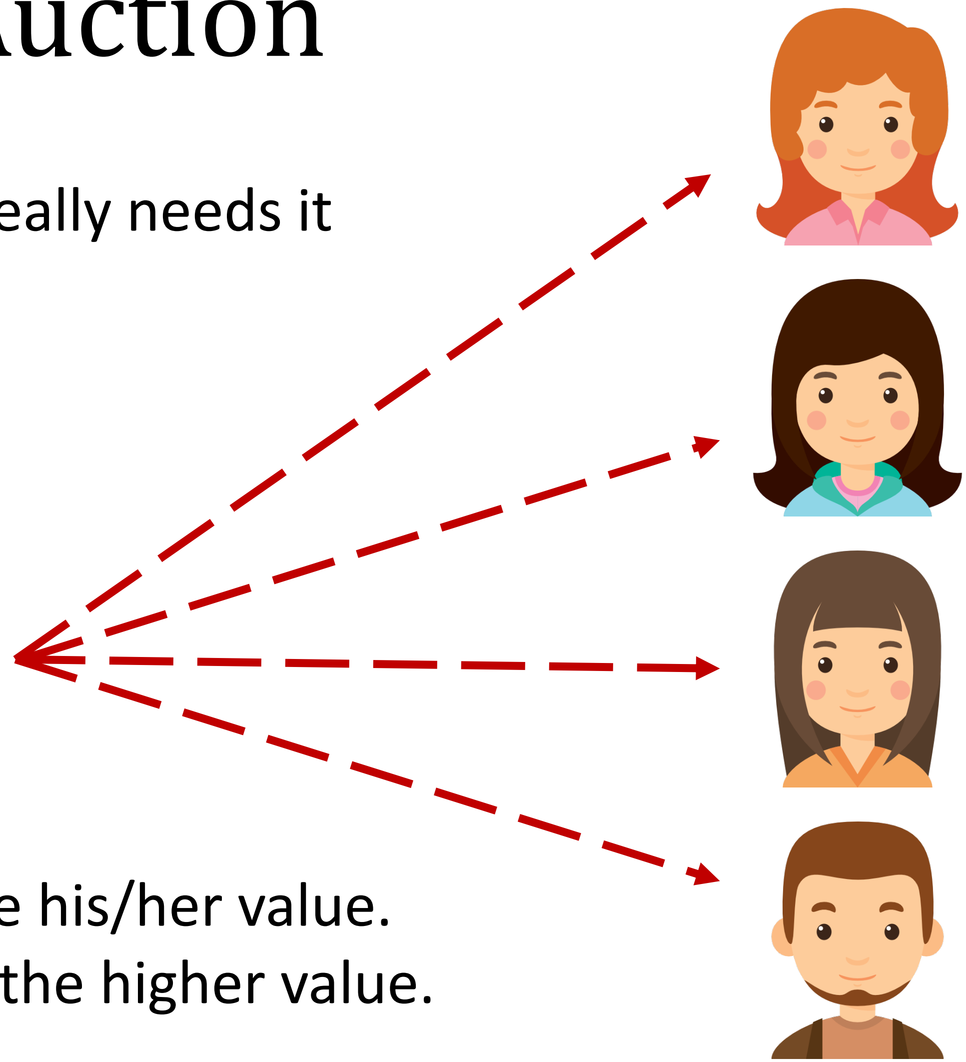

Rule 1: Each would tell me his/her value. I'll give it to the one with the higher value.

Image Courtesy: Freepik

Single-Item Auction

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

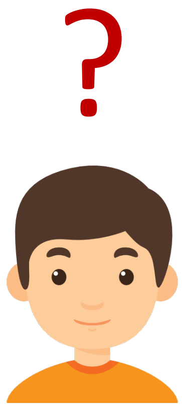

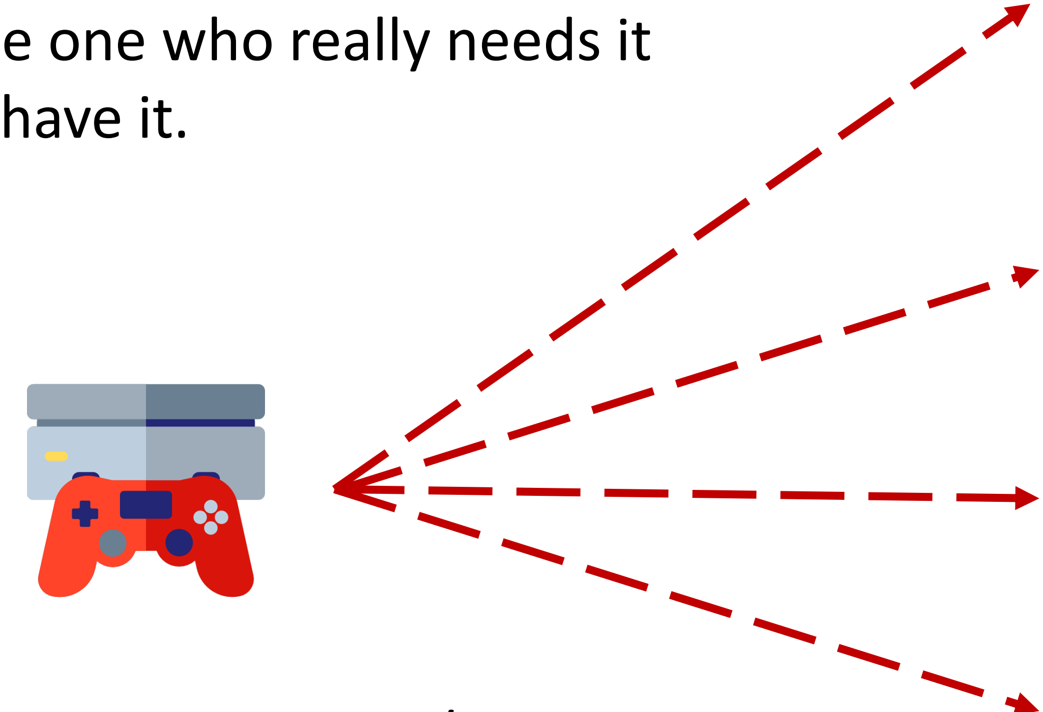

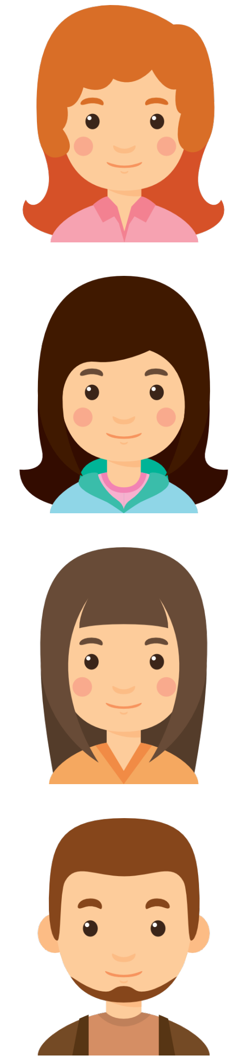

Rule 2: Each would tell me his/her value. I'll give it to the one with the higher value, but they have to pay me that value.

Image Courtesy: Freepik

Single-Item Auction

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

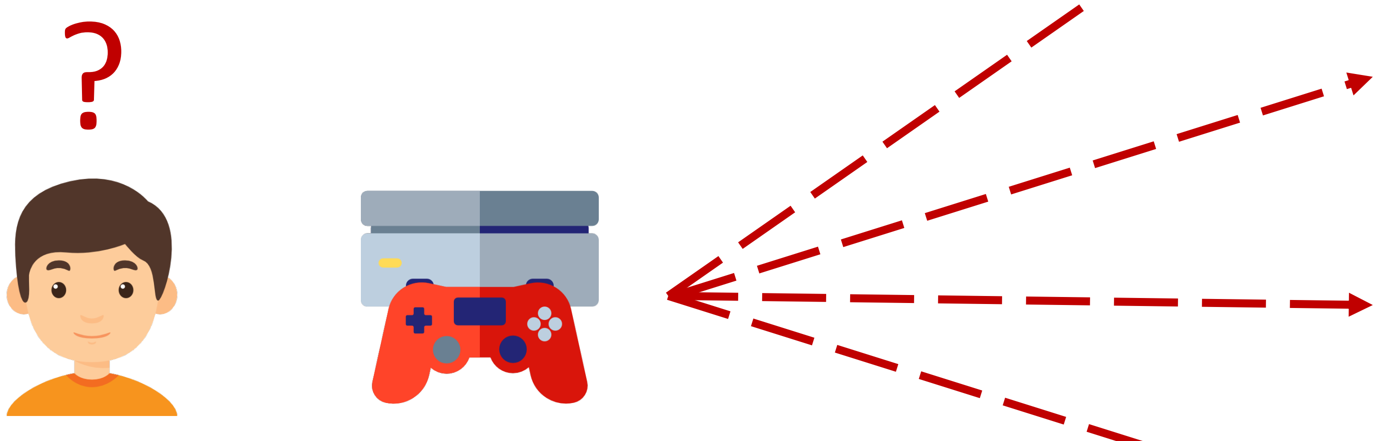

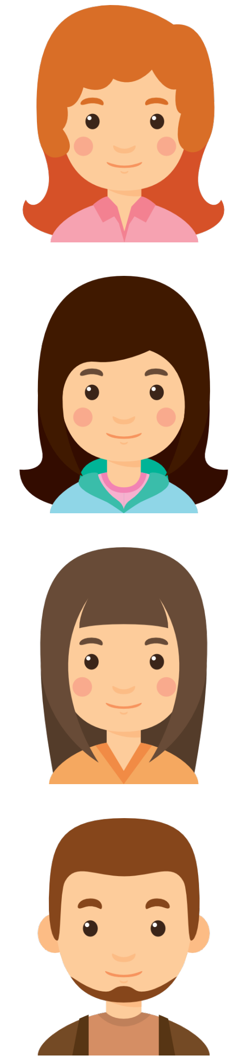

Implements the desired outcome.
But not truthfully.

Image Courtesy: Freepik

Single-Item Auction

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

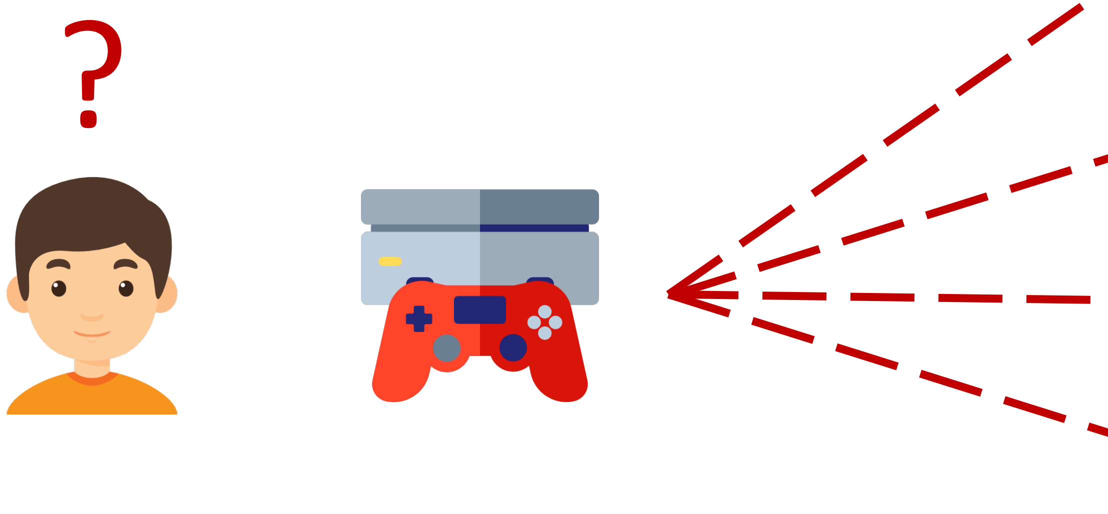

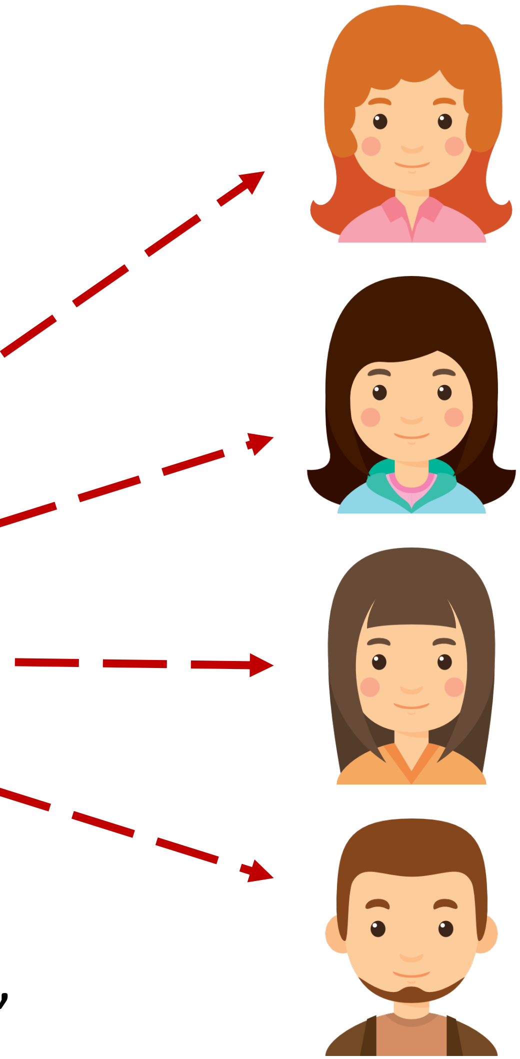

Rule 3: Each would tell me his/her value. I'll give it to the one with the highest value, and charge them the second highest value.

Image Courtesy: Freepik

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- Truthfulness/strategyproofness
  - For all agents i, all v_i, and all $\tilde{v}$,
$$u_i(v_i, \tilde{v}_{-i}) \geq u_i(\tilde{v}_i, \tilde{v}_{-i})$$
  - "Every agent is at least as happy reporting the truth as telling any lie, irrespective of what other agents report"
  - Almost same as telling the truth being a weakly dominant action
    - What's the difference?

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- Individual rationality
  - For all agents i and for all $\tilde{v}_{-i}$,
$$u_i(v_i, \tilde{v}_{-i}) \geq 0$$
  - "No agent should regret participating if she tells the truth."
  - Assumes that the utility from not participating is 0

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- No payments to agents
  - For all agents i and for all $\tilde{v}$,
$$p_i(\tilde{v}) \geq 0$$
  - "Agents pay the center. Not the other way around."
  - Common for auctions, but we may want the reverse in other settings

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- Welfare maximization
  - $f(\tilde{v})$ must be in $\text{argmax}_a \sum_i \tilde{v}_i(a)$
    - Important when making the users happy matters more than the immediate short-term revenue
    - Or think of the auctioneer as "agent $n+1$" with utility equal to the total payment received $\sum_i p_i(\tilde{v})$, and look at total utility

$$\left(\sum_i v_i\big(f(\tilde{v})\big) - p_i\big(f(\tilde{v})\big)\right) + \left(\sum_i p_i\big(f(\tilde{v})\big)\right) = \sum_i v_i\big(f(\tilde{v})\big)$$

Single-item Vickrey Auction

- Simplifying notation: v_i = value of agent i for the item
- $f(\tilde{v})$: give the item to agent $i^* \in \text{argmax}_i\ \tilde{v}_i$
- $p(\tilde{v})$: $p_{i^*} = \max_{j \neq i^*} \tilde{v}_j$, other agents pay nothing

Theorem:
Single-item Vickrey auction is strategyproof.

Proof sketch:

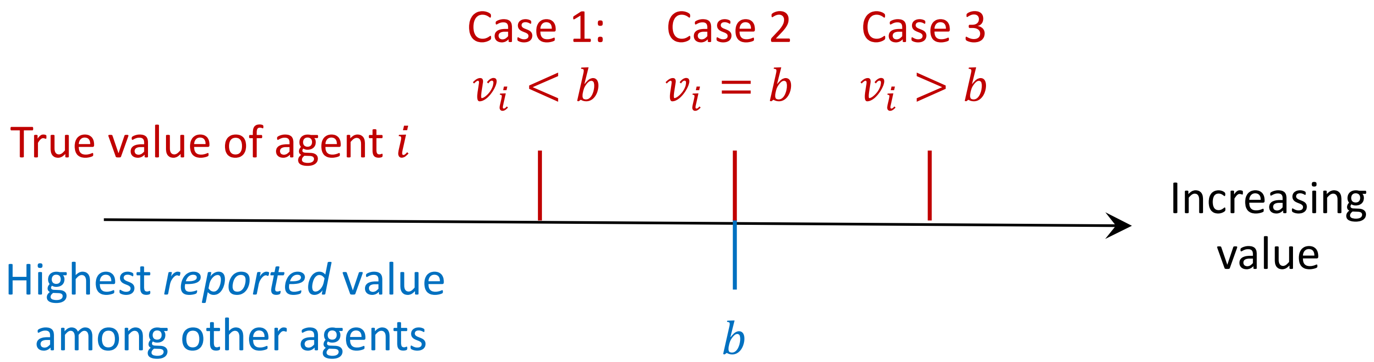

Vickrey Auction: Identical Items

- Two identical xboxes
  - Each agent i only wants one, has value v_i
  - Goal: give to the agents with the two highest values
- Attempt 1
  - To agent with highest value, charge 2^{nd} highest value.
  - To agent with 2^{nd} highest value, charge 3^{rd} highest value.
- Attempt 2
  - To agents with highest and 2^{nd} highest values, charge the 3^{rd} highest value.
- Question: Which attempt(s) would be strategyproof?
  - Both, 1, 2, None?

VCG Auction

- Recall the general setup:
  - A = set of outcomes, v_i = valuation of agent i, $\tilde{v}_i$ = what agent i reports, f chooses the outcome, p decides payments
- VCG (Vickrey-Clarke-Groves Auction)
  - $f(\tilde{v}) = a^* \in \text{argmax}_{a \in A} \sum_i \tilde{v}_i(a)$
  - $p_i(\tilde{v}) = \left[\max_a \sum_{j \neq i} \tilde{v}_j(a)\right] - \left[\sum_{j \neq i} \tilde{v}_j(a^*)\right]$

Maximize welfare

i's payment = welfare that others lost due to presence of i

A Note About Payments

- $p_i(\tilde{v}) = \left[\max_a \sum_{j \neq i} \tilde{v}_j(a)\right] - \left[\sum_{j \neq i} \tilde{v}_j(a^*)\right]$

- In the first term…
  - Maximum is taken over alternatives that are feasible when i does not participate.
  - Agent i cannot affect this term, so can ignore in calculating incentives.
  - Could be replaced with any function $h_i(\tilde{v}_{-i})$
    - This specific function has advantages (we'll see)

VCG: Simple Example

- Suppose each agent has a value XBox and a value for PS4.
- Their value for $\{XBox, PS4\}$ is the max of their two values.

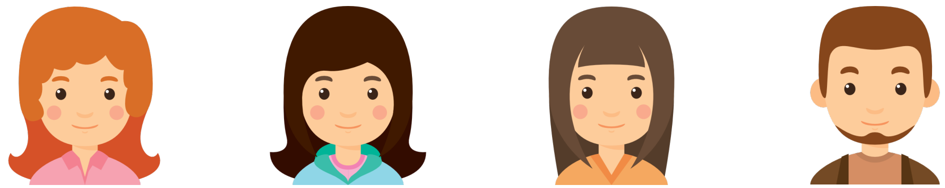

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Q: Who gets the xbox and who gets the PS4?

Q: How much do they pay?

VCG: Simple Example

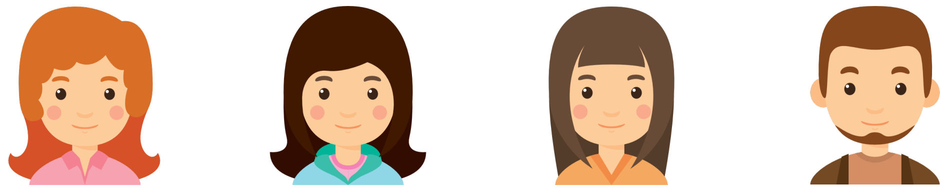

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Allocation:

- A4 gets XBox, A3 gets PS4
- Achieves maximum welfare of $7 + 6 = 13$

VCG: Simple Example

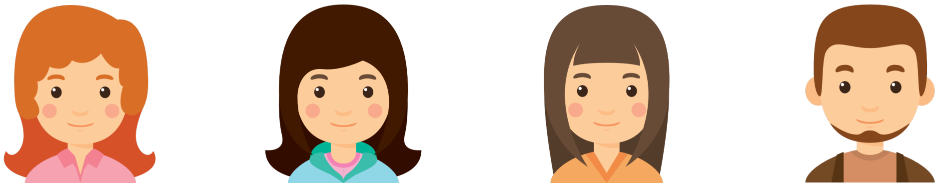

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Payments:

- Zero payments charged to A1 and A2
  - “Deleting” either does not change the outcome/payments for others
- Can also be seen by individual rationality

VCG: Simple Example

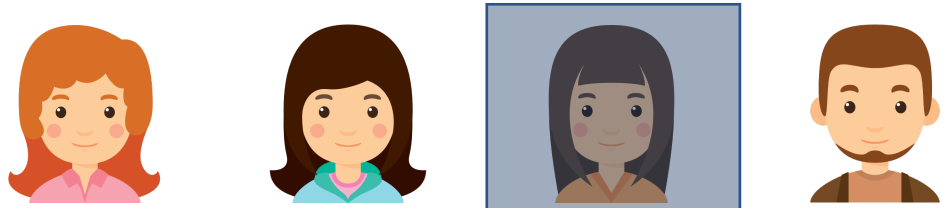

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Payments:

- Payment charged to A3 $= 11 - 7 = 4$
  - Max welfare to others if A3 absent: $7 + 4 = 11$
    - Give XBox to A4 and PS4 to A1
  - Welfare to others if A3 present: 7

VCG: Simple Example

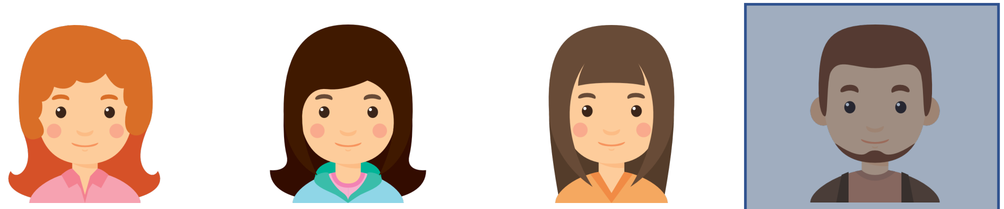

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Payments:

- Payment charged to A4 $= 12 - 6 = 6$
  - Max welfare to others if A4 absent: $8 + 4 = 12$
    - Give XBox to A3 and PS4 to A1
  - Welfare to others if A4 present: 6

VCG: Simple Example

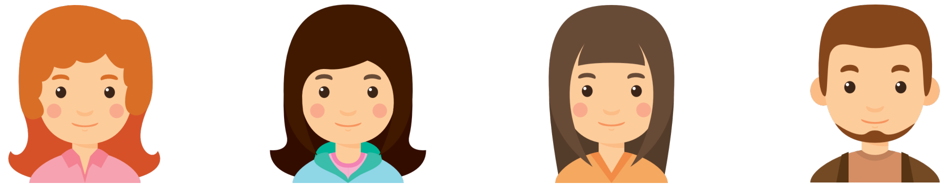

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Final Outcome:

- Allocation: A3 gets PS4, A4 gets XBox
- Payments: A3 pays 4, A4 pays 6
- Net utilities: A3 gets $6 - 4 = 2$, A4 gets $7 - 6 = 1$