CSC2556

Lecture 9

Mechanism Design with Money (VCG)

Announcements

- Class moving online starting next week
  - Starting 3/20, and for the remaining semester, our class is moving online
  - That means:
    - Lectures and project presentations will be online
      - Exact details TBD; I'll send out the instructions next week
    - Office hours will be through Skype
    - Homework and project reports will still be through MarkUs
- Mid-project (virtual) check-in
  - I'll send out a sign-up sheet during the weekend
    - Can sign up for a 30-minute slot to chat about progress in your project (voluntary)

VCG

Mathematical Setup

- A set of outcomes A
  - A might depend on which agents are participating.
- Each agent i has a private valuation $v_i : A \to \mathbb{R}$
- Auctions:
  - A has a nice structure.
    - Selling one item to n buyers $= n$ outcomes ("give to i")
    - Selling m items to n buyers $= n^m$ outcomes
  - Agents only care about which items *they* receive
    - A_i = bundle of items allocated to agent i
    - Use $v_i(A_i)$ instead of $v_i(A)$ for notational simplicity
  - But for now, we'll look at the general setup.

Mathematical Setup

- Agent i might misreport: report $\tilde{v}_i$ instead of v_i

- Mechanism: (f, p)
  - Input: reported valuations $\tilde{v} = (\tilde{v}_1, \dots, \tilde{v}_n)$
  - $f(\tilde{v}) \in A$ decides what outcome is implemented
  - $p(\tilde{v}) = (p_1, \dots, p_n)$ decides how much each agent pays
    - Note that each p_i is a function of all reported valuations

- Utility to agent i : $u_i(\tilde{v}) = v_i\big(f(\tilde{v})\big) - p_i(\tilde{v})$
  - "Quasi-linear utilities"

Mathematical Setup

- Our goal is to design the mechanism (f, p)
  - f is called the social choice function
  - p is called the payment scheme
  - We want to several things from our mechanism

- Truthfulness/strategyproofness
  - For all agents i, all v_i, and all $\tilde{v}$,
$$u_i(v_i, \tilde{v}_{-i}) \geq u_i(\tilde{v}_i, \tilde{v}_{-i})$$
  - An agent is at least as happy reporting the truth as telling any lie, irrespective of what other agents report

Mathematical Setup

- Our goal is to design the mechanism (f, p)
  - f is called the social choice function
  - p is called the payment scheme
  - We want to several things from our mechanism

- Individual rationality
  - For all agents i and for all $\tilde{v}_{-i}$,
$$u_i(v_i, \tilde{v}_{-i}) \geq 0$$
  - An agent doesn't regret participating if she tells the truth.

Mathematical Setup

- Our goal is to design the mechanism (f, p)
  - f is called the social choice function
  - p is called the payment scheme
  - We want to several things from our mechanism

- 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.

Mathematical Setup

- Our goal is to design the mechanism (f, p)
  - f is called the social choice function
  - p is called the payment scheme
  - We want to several things from our mechanism

- Welfare maximization
  - Maximize $\sum_i v_i\left(f(\tilde{v})\right)$
    - In many contexts, payments are less important (e.g. ad auctions)
    - Or think of the auctioneer as another agent with utility $\sum_i p_i(\tilde{v})$
      - Then, the total utility of all agents (including the auctioneer) is precisely the objective written above

Single-Item Auction

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

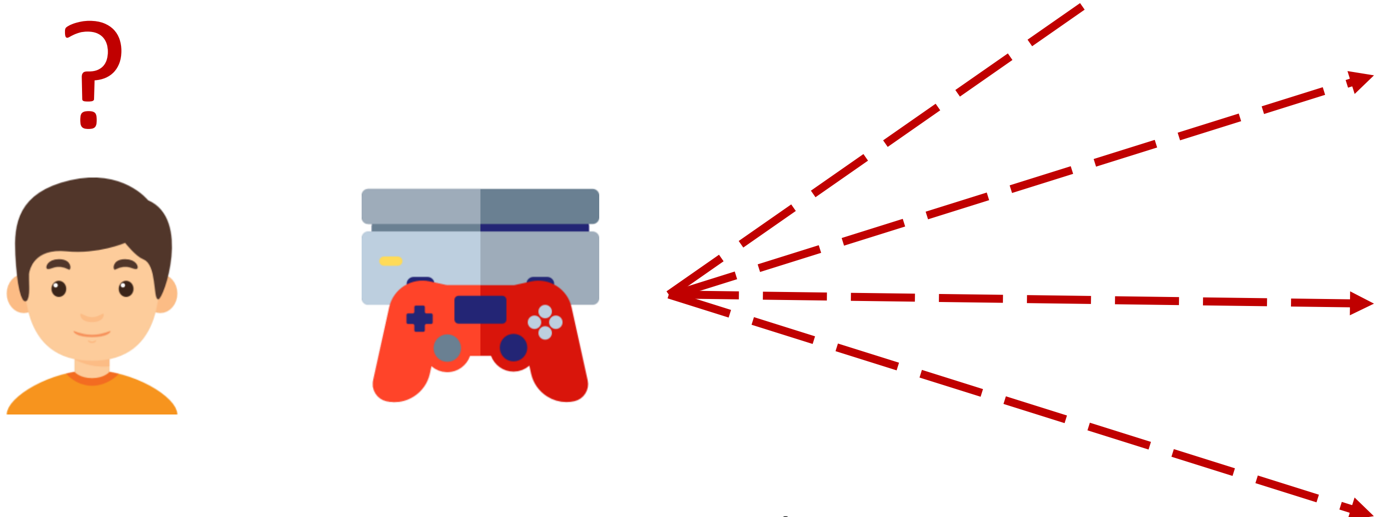

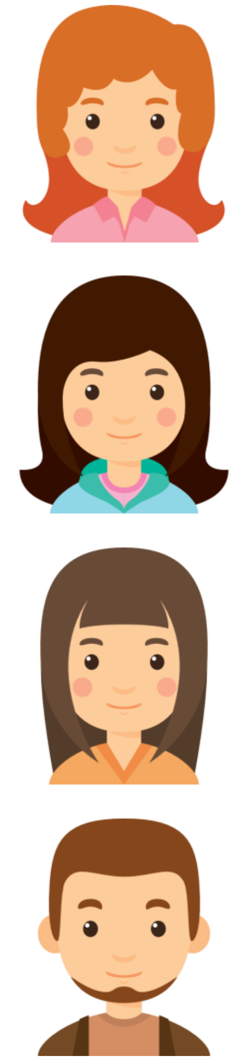

Rule 1: Each would tell me his/her value.
I'll give it to the one with the higher value.

Image Courtesy: Freepik

Single-Item Auction

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

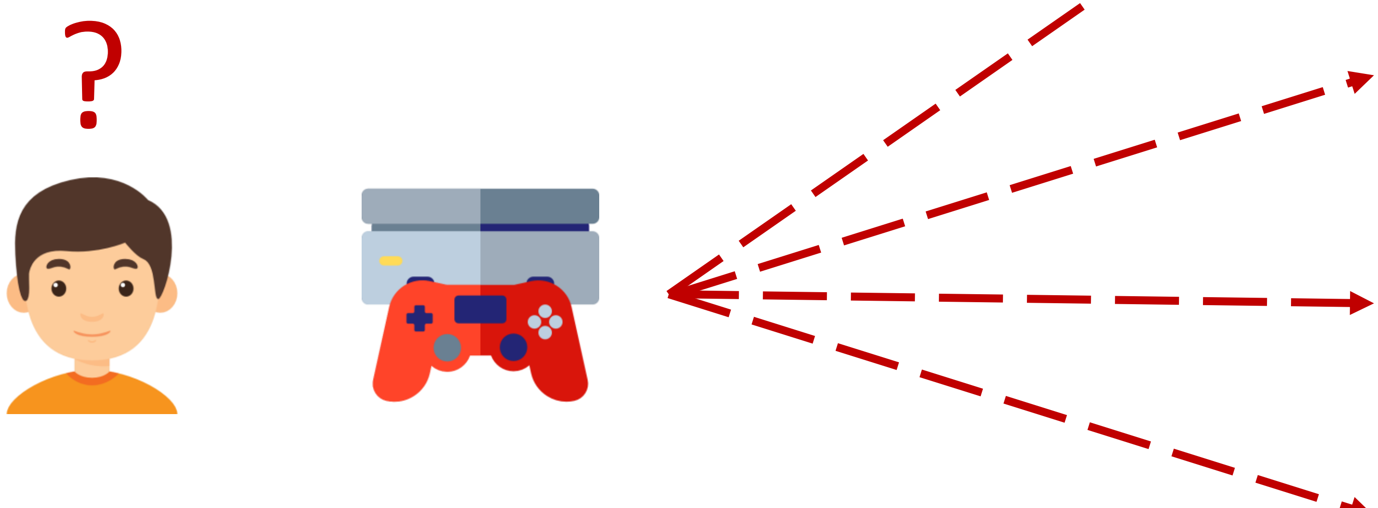

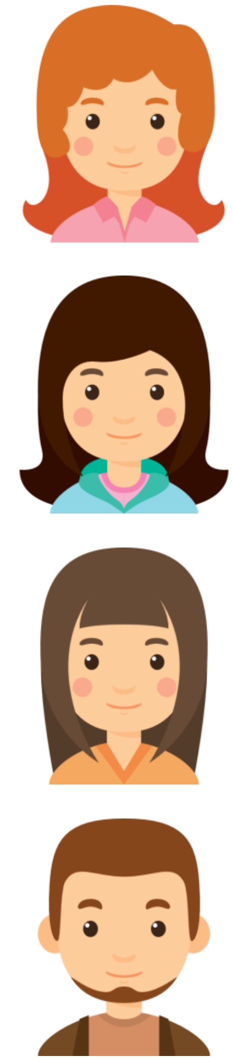

Rule 2: Each would tell me his/her value. I'll give it to the one with the higher value, but they have to pay me that value.

Image Courtesy: Freepik

Single-Item Auction

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

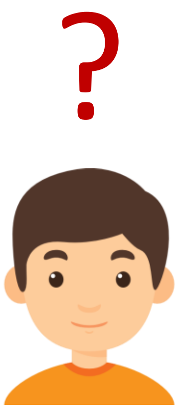

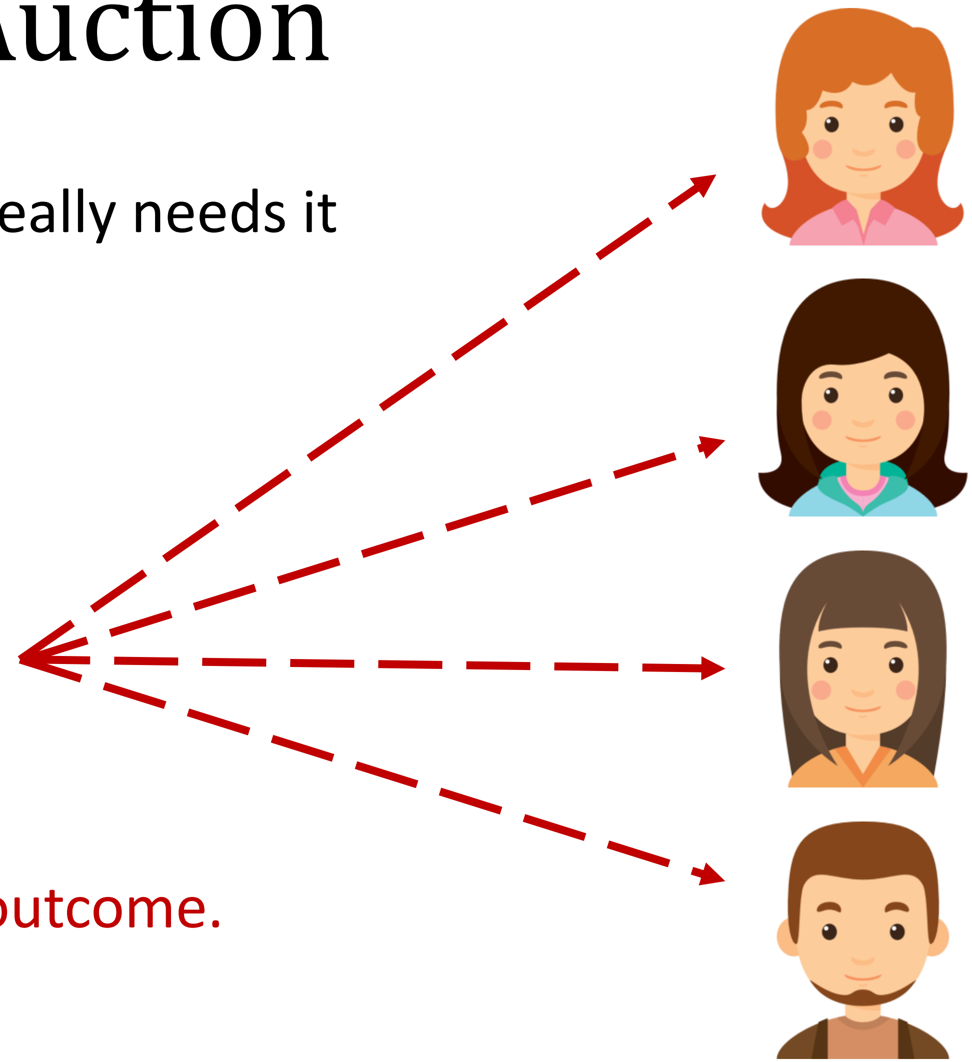

Implements the desired outcome.
But not truthfully.

Image Courtesy: Freepik

Single-Item Auction

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

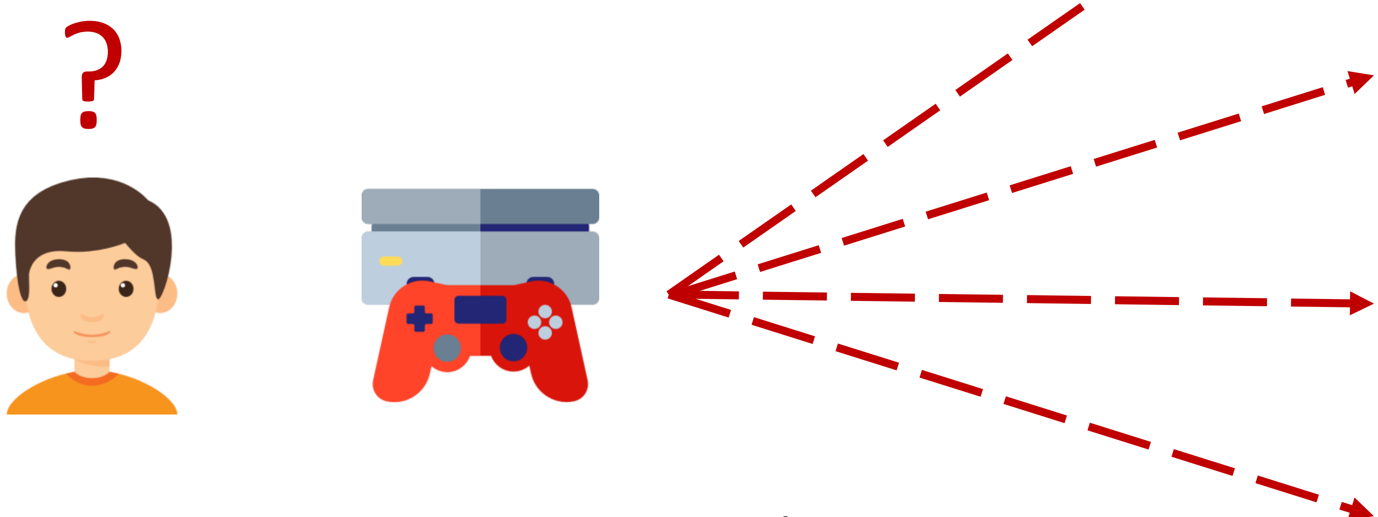

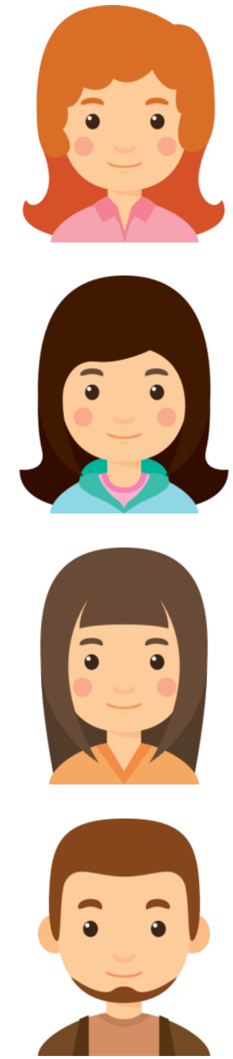

Rule 3: Each would tell me his/her value. I'll give it to the one with the highest value, and charge them the second highest value.

Image Courtesy: Freepik

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:

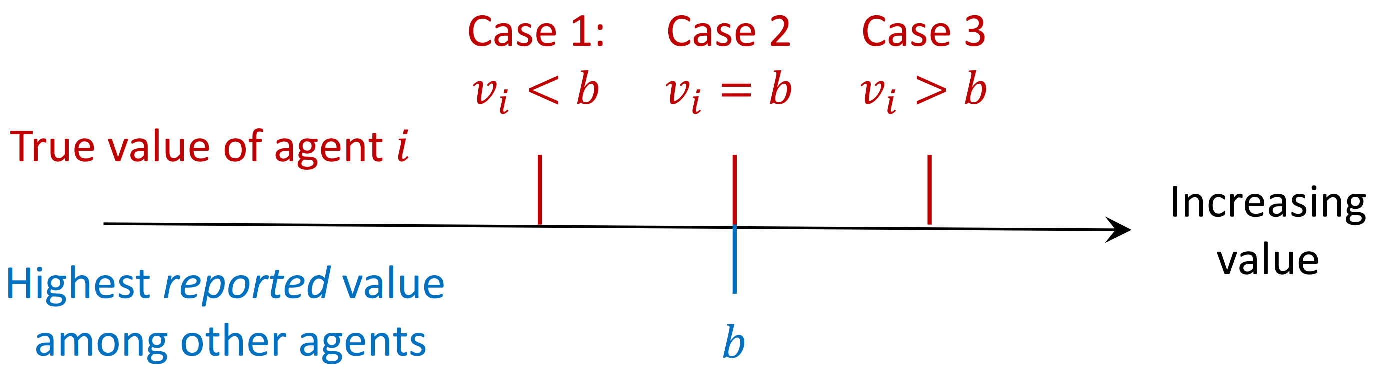

Vickrey Auction: Identical Items

- 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 2nd highest value.
  - To agent with 2nd highest value, charge 3rd highest value.
- Attempt 2
  - To agents with highest and 2nd highest values, charge the 3rd 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)
  - $f(\tilde{v}) = a^* \in \text{argmax}_{a \in A} \sum_i \tilde{v}_i(a)$

Maximize welfare

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

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)

Properties of VCG Auction

- Strategyproofness:
  - Suppose agents other than i report $\tilde{v}_{-i}$.
  - Agent i reports $\tilde{v}_i \Rightarrow$ outcome chosen is $f(\tilde{v}) = a$
  - Utility to agent i = $v_i(a) - \left(\blacksquare - \sum_{j \neq i} \tilde{v}_j(a)\right)$

    Term that agent i cannot affect

  - Agent i wants a to maximize $v_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
  - f chooses a to maximize $\tilde{v}_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
  - Hence, agent i is best off reporting $\tilde{v}_i = v_i$
    - f chooses a that maximizes the utility to agent i

Properties of VCG Auction

- Individual rationality:
  - $a^* \in \text{argmax}_{a \in A}\, v_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
  - $\tilde{a} \in \text{argmax}_{a \in A}\, \sum_{j \neq i} \tilde{v}_j(a)$

$$
\begin{aligned}
& u_i(v_i, \tilde{v}_{-i}) \\
& = v_i(a^*) - \left(\sum_{j \neq i} \tilde{v}_j(\tilde{a}) - \sum_{j \neq i} \tilde{v}_j(a^*) \right) \\
& = \left[v_i(a^*) + \sum_{j \neq i} \tilde{v}_j(a^*) \right] - \left[\sum_{j \neq i} \tilde{v}_j(\tilde{a}) \right] \\
& = \text{Max welfare to all agents} \\
& \qquad - \text{max welfare to others when } i \text{ is absent} \\
& \geq 0
\end{aligned}
$$

Properties of VCG Auction

- No payments to agents:
  - Suppose the agents report $\tilde{v}$
  - $a^* \in \text{argmax}_{a \in A} \sum_j \tilde{v}_j(a)$
  - $\tilde{a} \in \text{argmax}_{a \in A} \sum_{j \neq i} \tilde{v}_j(a)$

$$
\begin{aligned}
& p_i(\tilde{v}) \\
&= \sum_{j \neq i} \tilde{v}_j(\tilde{a}) - \sum_{j \neq i} \tilde{v}_j(a^*) \\
&= \text{Max welfare to others when } i \text{ is absent} \\
&\qquad - \text{welfare to others when } i \text{ is present} \\
&\geq 0
\end{aligned}
$$

Properties of VCG Auction

- Welfare maximization:
  - By definition, since f chooses the outcome maximizing the sum of reported values

- Informal result:
  - Under minimal assumptions, VCG is the unique auction satisfying these properties.

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.

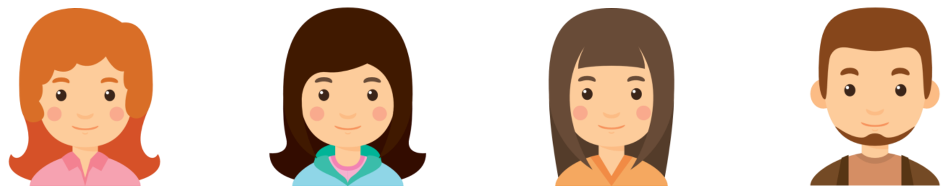

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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

Q: How much do they pay?

VCG: Simple Example

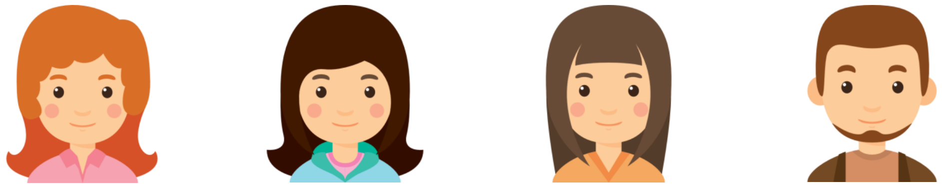

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Allocation:

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

VCG: Simple Example

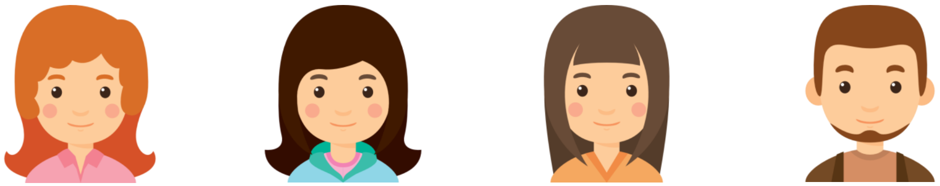

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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

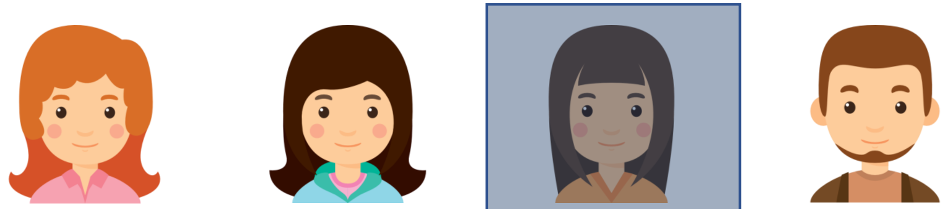

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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

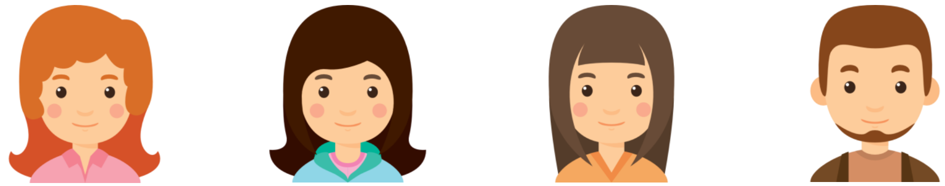

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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 $7 - 6 = 1$

Problems with VCG

- Difficult to understand
  - Must reason about what would maximize others' welfare
- Possibly low revenue
  - [Bulow-Klemperer 96]: With i.i.d. valuations, $\mathbb{E}$[VCG revenue, n+1 agents] $\geq$ $\mathbb{E}$[OPT revenue, n agents]
- Often NP-hard to implement
  - Even computing the welfare maximizing allocation may be computationally difficult
- …

Single-Minded Bidders

- Allocate a set S of m items

- Each agent i is described by (v_i, S_i)
  - Gets value v_i if she receives all items in $S_i \subseteq S$ (and possibly some other items)
  - Gets value 0 if she doesn't receive even one item in S_i
  - "Single-minded"

- Welfare-maximizing allocation:
  - Find a subset of players with the highest total value such that their desired sets are disjoint

Single-Minded Bidders

- Reduction to the Weighted Independent Set (WIS) problem in graphs
  - NP-hard
  - No $O(m^{0.5-\epsilon})$ approximation (unless $NP \subseteq ZPP$)

- $\sqrt{m}$-approximation through a simple greedy algorithm *in a strategyproof way*

Greedy Algorithm

- Input: (v_i, S_i) for each agent i
- Output: Agents with mutually independent S_i

- Greedy Algorithm:
  - Sort the agents in a specific order (we'll see).
  - Relabel them as $1,2, \ldots, n$ in this order.
  - $W \leftarrow \emptyset$
  - For $i = 1, \ldots, n$:
    - If $S_i \cap S_j = \emptyset$ for every $j \in W$, then $W \leftarrow W \cup \{i\}$
  - Give agents in W their desired items.

Greedy Algorithm

- Sort by what?

- We want to satisfy agents with higher values.
  - $v_1 \geq v_2 \geq \cdots \geq v_n \Rightarrow m$-approximation ☹

- But we don't want to exhaust too many items.
  - $\frac{v_1}{|S_1|} \geq \frac{v_2}{|S_2|} \geq \cdots \frac{v_n}{|S_n|} \Rightarrow m$-approximation ☹

- $\sqrt{m}$-approximation : $\frac{v_1}{\sqrt{|S_1|}} \geq \frac{v_2}{\sqrt{|S_2|}} \geq \cdots \frac{v_n}{\sqrt{|S_n|}}$?

[Lehmann et al. 2011]

Proof of Approximation

- Definitions
  - OPT = Agents satisfied by the optimal algorithm
  - W = Agents satisfied by the greedy algorithm
  - For $i \in W$,
$$OPT_i = \{j \in OPT, j \geq i : S_i \cap S_j \neq \emptyset\}$$
- Claim 1: $OPT \subseteq \bigcup_{i \in W} OPT_i$
- Claim 2: It is enough to show that $\forall i \in W$
$$\sqrt{m} \cdot v_i \geq \Sigma_{j \in OPT_i} v_j$$
- Observation: For $j \in OPT_i$, $v_j \leq v_i \cdot \frac{\sqrt{|S_j|}}{\sqrt{|S_i|}}$

Proof of Approximation

- Summing over all $j \in OPT_i$:

$$\Sigma_{j \in OPT_i} v_j \leq \frac{v_i}{\sqrt{|S_i|}} \cdot \Sigma_{j \in OPT_i} \sqrt{|S_j|}$$

- Using Cauchy-Schwarz ($\Sigma_i x_i y_i \leq \sqrt{\Sigma_i x_i^2} \cdot \sqrt{\Sigma_i y_i^2}$)

$$\Sigma_{j \in OPT_i} \sqrt{|S_j| \cdot 1} \leq \sqrt{|OPT_i|} \cdot \sqrt{\Sigma_{j \in OPT_i} |S_j|}$$

$$\leq \sqrt{|S_i|} \cdot \sqrt{m}$$

Strategyproofness

- Agent i pays $p_i = v_{j^*} \cdot \sqrt{\frac{|S_i|}{|S_{j^*}|}}$
  - j^* is the smallest index j such that j is currently not selected by greedy but would be selected if we remove (v_i, S_i) from the system
  - Exercise: Show that we must have $j^* > i$
  - Exercise: Show that $S_i \cap S_{j^*} \neq \emptyset$
  - Another interpretation: p_i = lowest value i can report and still win

Strategyproofness

- Critical payment
  - Charge each agent the lowest value they can report and still win
- Monotonic allocation
  - If agent i wins when reporting (v_i, S_i), she must win when reporting $v_i' \geq v_i$ and $S_i' \subseteq S_i$.
  - Greedy allocation rule satisfies this.
- Theorem: Critical payment + monotonic allocation rule imply strategyproofness.

Moral

- VCG can sometimes be too difficult to implement
  - May look into approximately maximizing welfare
  - As long as the allocation rule is monotone, we can charge critical payments to achieve strategyproofness
  - Note: approximation is needed for computational reasons
- Use of approximation
  - Note that here we used approximation to circumvent computational hardness, not to achieve strategyproofness
  - In mechanism design without money, we needed approximation even to just achieve strategyproofness