CSC2556

Lecture 8

Mechanism Design with Money: VCG

Announcements

- Mid-project Check-in:
  - Sent out a sign-up sheet.
  - If you think it would help, sign up for a 30-minute slot and we can chat about your project.
- Presentations:
  - We'll have presentations in the last 1.5 lectures with about 20 minutes per group (17 minutes of presentation followed by 3 minutes of class discussion).
- Reports: due April 15
  - 4-5 pages
  - Introduction, related work, model, results, future work

Framework

- Set N of n agents
- Set A of m alternatives
- Valuations $v = (v_i)_{i \in N}$
  - Agent i's valuation: $v_i: A \to \mathbb{R}$
- Mechanism $M = (f, p)$
  - Social Choice Function: $f(v) \in A$ is implemented
  - Payment Vector: Agent i pays $p_i(v)$

Framework

- Quasi-linear utilities: $v_i\big(f(v)\big) - p_i(v)$
- Goal 1: Social Welfare Maximization
  - Maximize $\sum_i v_i\big(f(v)\big)$
  - Can think of welfare with auctioneer. Also important to generate high-quality ads in ad auctions.
- Goal 2: Revenue Maximization (we'll skip this)
  - Maximize $\sum_i p_i(v)$
- Individual Rationality (IR)
  - Non-negative utilities: $v_i\big(f(v)\big) - p_i(v) \geq 0, \forall i \in N$
  - Bounds the revenue in goal 2.

Framework

- Difficulty:
  - Agents may report incorrect valuations $\tilde{v} = (\tilde{v}_i)_{i \in N}$
  - Agent i, given the reports of other agents $\tilde{v}_{-i}$, wants to maximize her own utility $v_i\big(f(\tilde{v}_i, \tilde{v}_{-i})\big) - p_i(\tilde{v}_i, \tilde{v}_{-i})$

- Strategyproofness (SP)
  - Each agent i maximizes her utility by reporting her true valuation v_i, regardless of what other agents report.

    $v_i \in \text{argmax}_{\tilde{v}_i}\, v_i\big(f(\tilde{v}_i, \tilde{v}_{-i})\big) - p_i(\tilde{v}_i, \tilde{v}_{-i}), \forall i, \tilde{v}_{-i}$
  - Achieving SP is why we'll need to charge payments in Goal 1.

Auctions

- Sell a set of goods to a set of agents
  - Similar to fair division, but now with payments
  - Alternative $a \rightarrow$ allocation A
  - Standard assumption:
    - Agent i's value only depends on A_i
    - Instead of $v_i(a)$, we use $v_i(A_i)$

- Single-item Auction
  - Alternative a_i : "agent i gets the item"
  - $v_i(a_i) \rightarrow v_i$ (shorthand), $v_i(a_j) = 0, \forall i \neq j$

Single-Item Auction

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

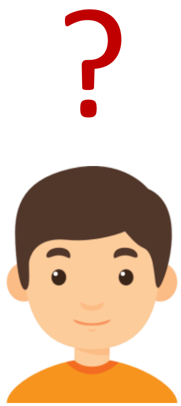

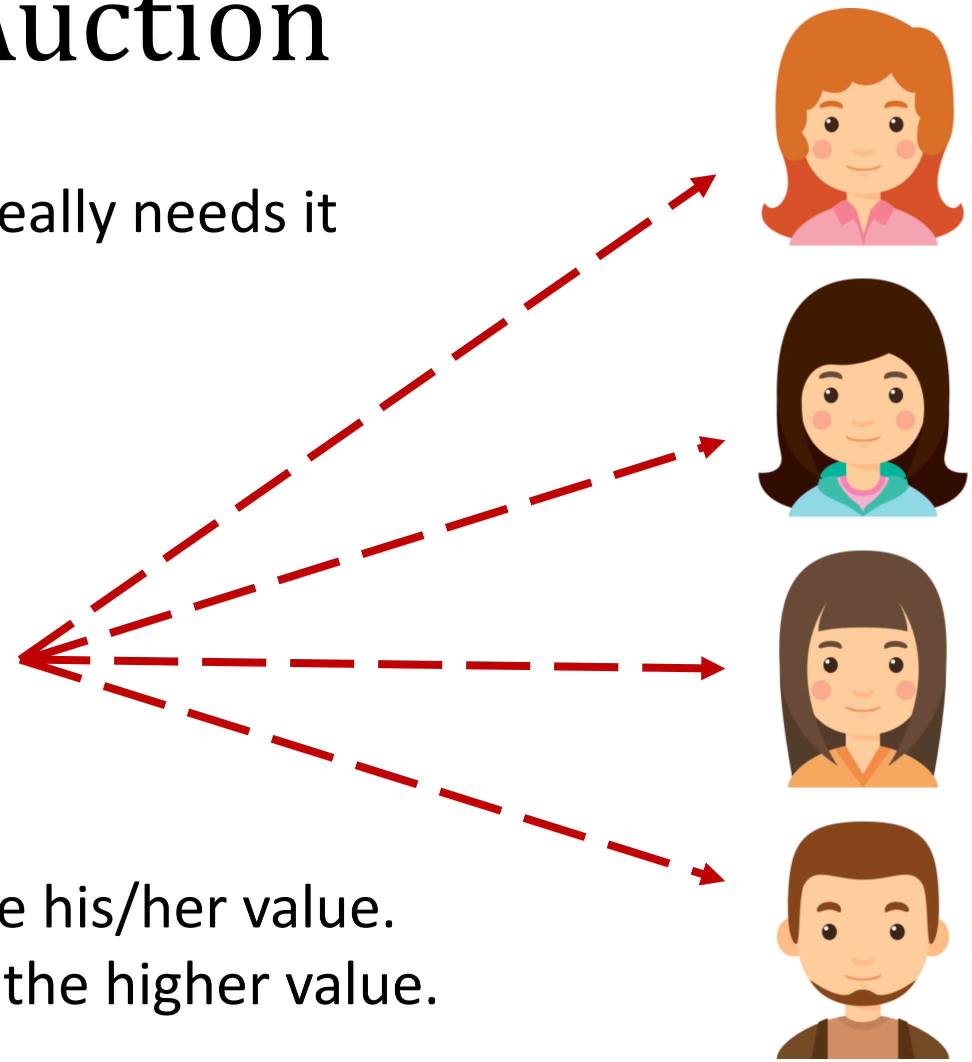

Rule 1: Each would tell me his/her value. I'll give it to the one with the higher value.

Image Courtesy: Freepik

Single-Item Auction

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

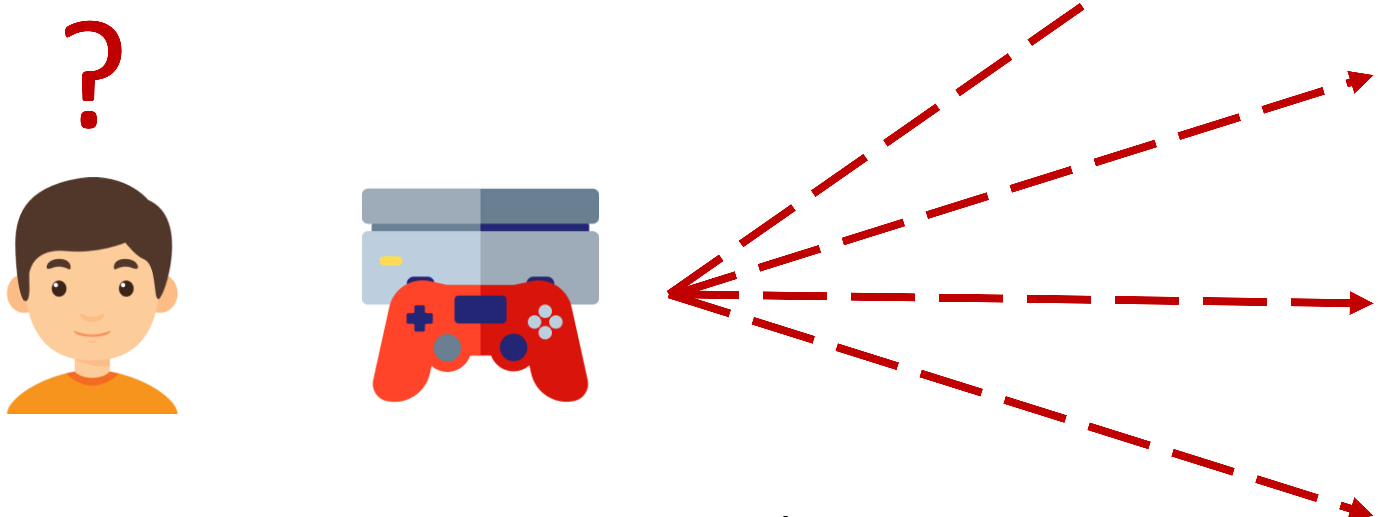

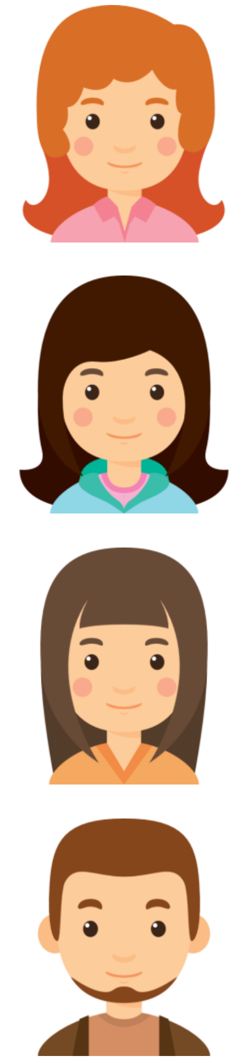

Rule 2: Each would tell me his/her value. I'll give it to the one with the higher value, but they have to pay me that value.

Image Courtesy: Freepik

Single-Item Auction

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

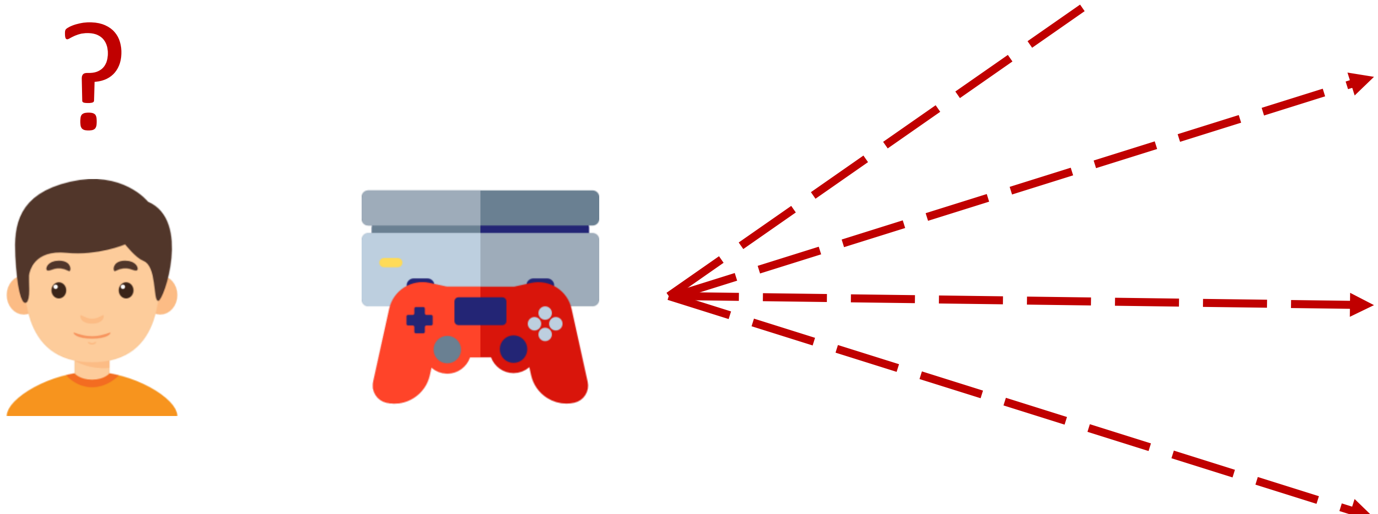

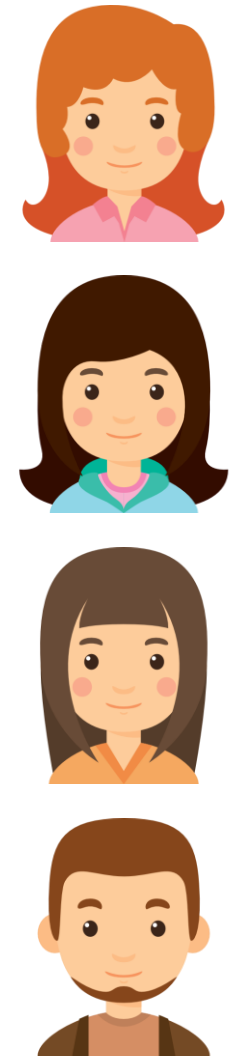

Implements the desired outcome.
But not in a strategyproof way.

Image Courtesy: Freepik

Single-Item Auction

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

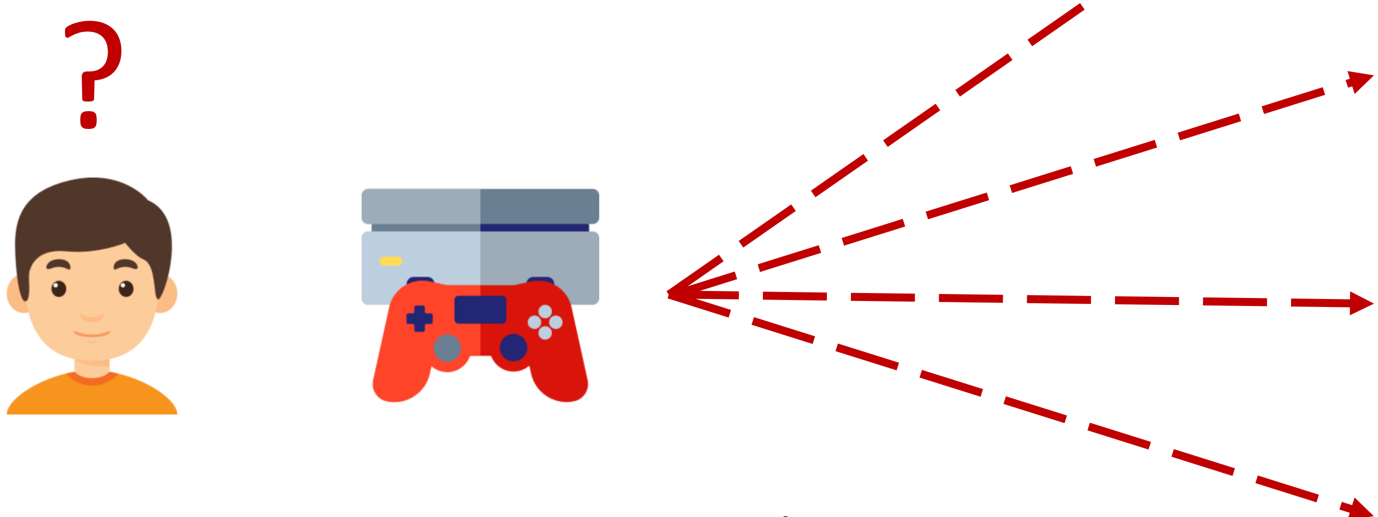

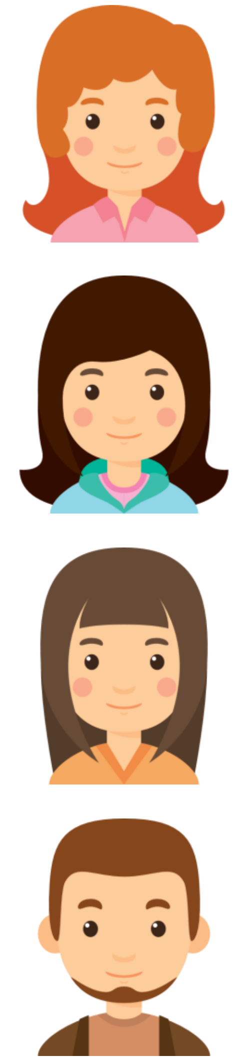

Rule 3: Each would tell me his/her value. I'll give it to the one with the highest value, and charge them the second highest value.

Image Courtesy: Freepik

VCG: Single-Item

- f : Give the item to agent $i^* \in \text{argmax}_i\, v_i$
- p : $p_{i^*} = \max_{j \neq i^*} v_j$, other agents pay nothing

Theorem:
VCG for a single item is strategyproof.

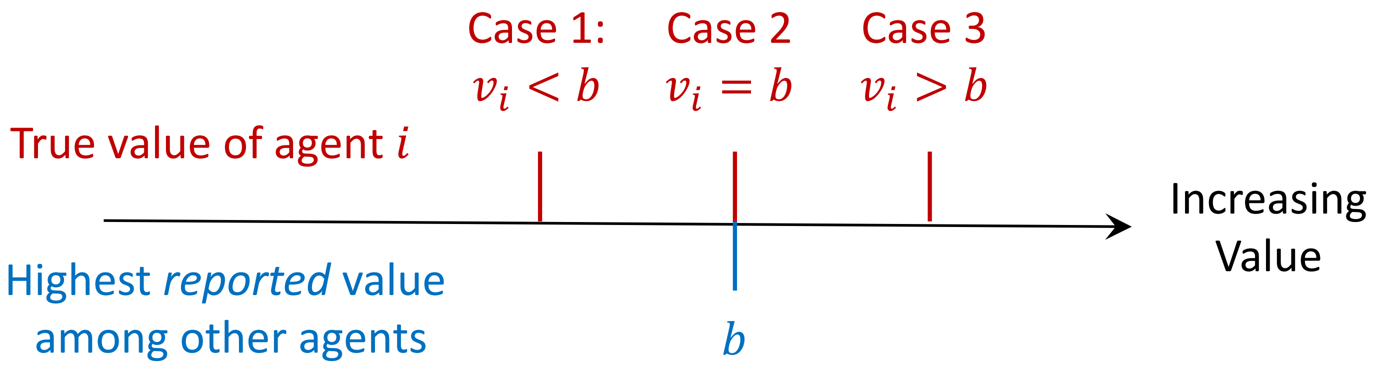

VCG: Identical Items

- 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:
  - Highest value → pay 2^{nd} highest value
  - 2^{nd} highest value → pay 3^{rd} highest value
- Attempt 2:
  - {Highest value, 2^{nd} highest value} → pay 3^{rd} highest value
- Question: Which would be strategyproof?

Vickrey Auction: General Case

- For the general case with arbitrary alternatives

- Vickrey Auction
  - $f(v) = \text{argmax}_{a \in A} \sum_i v_i(a)$
  - $p_i(v) = -\sum_{j \neq i} v_j\big(f(v)\big)$

Maximize social welfare

Pay (not charge!) to each agent the total value to others

- Why is this SP?
  - Suppose agent $j \neq i$ reports $\tilde{v}_j$
  - Utility to agent i when reporting $\tilde{v}_i$
    - $v_i(a) - \big(-\sum_{j \neq i} \tilde{v}_j(a)\big) = v_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
    - Mechanism chooses a to maximize $\tilde{v}_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
    - Utility maximized when reporting $\tilde{v}_i = v_i$

Vickrey Auction

- This achieves social welfare maximization and individual rationality (IR)

- But: To give away my single xbox, I need to pay each friend who doesn't get it the value of the friend who gets it (I'm not that rich!)

- Additional property:
  - Agents pay the principal: $p_i(v) \geq 0$

Idea

- Vickrey auction
  - $f(v) = \text{argmax}_{a \in A} \sum_i v_i(a)$
  - $p_i(v) = -\sum_{j \neq i} v_j\big(f(v)\big)$

- A slight modification
  - $f(v) = \text{argmax}_{a \in A} \sum_i v_i(a)$
  - $p_i(v) = h_i(v_{-i}) - \sum_{j \neq i} v_j\big(f(v)\big)$

- Still truthful. Agent i has no control over his additional payment $h_i(v_{-i})$

VCG

- Clarke's pivot rule
  - $h_i(v_{-i}) = \max_a \sum_{j \neq i} v_j(a)$
  - Maximum welfare to others if agent i wasn't there

- VCG (Vickrey-Clarke-Groves Auction)
  - $f(v) = a^* = \text{argmax}_{a \in A} \sum_i v_i(a)$
  - $p_i(v) = \left[\max_a \sum_{j \neq i} v_j(a)\right] - \left[\sum_{j \neq i} v_j(a^*)\right]$

- Payment charged to agent i = harm imposed on the welfare of others by i's presence

VCG

- $f(v) = a^* = \text{argmax}_{a \in A} \sum_i v_i(a)$
- $p_i(v) = \left[\max_a \sum_{j \neq i} v_j(a)\right] - \left[\sum_{j \neq i} v_j(a^*)\right]$

- We already saw that this is strategyproof.

- We also have $p_i(v) \geq 0$. (Why?)

- We maintain IR: $p_i(v) \leq v_i(a^*)$. (Why?)

VCG: Simple Example

- Let's go back to giving away an xbox and a ps4.

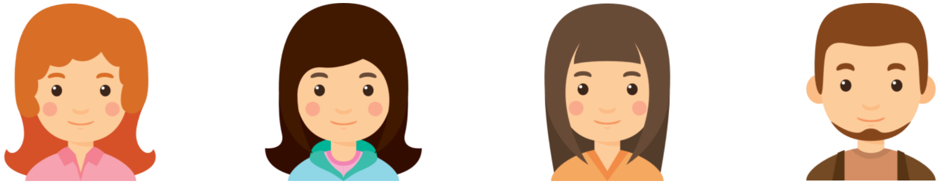

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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

Q: How much do they pay?

VCG: Simple Example

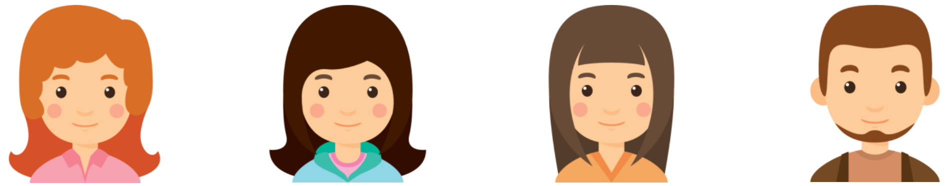

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Allocation:

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

VCG: Simple Example

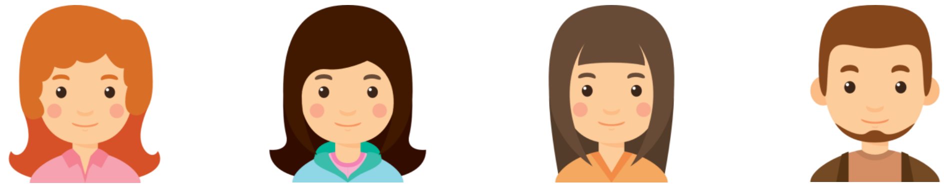

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Payments:

- Zero payments charged to A1 and A2
- "Deleting" either of them does not change the outcome or payments for others
- Can also be seen by individual rationality

VCG: Simple Example

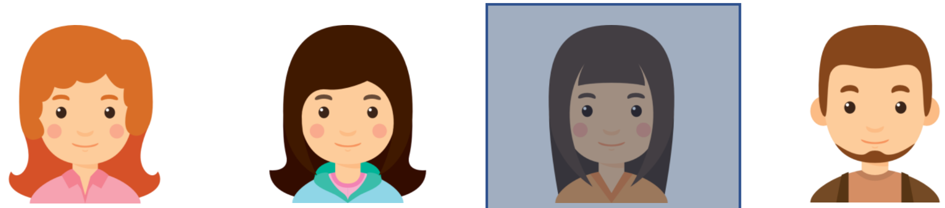

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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

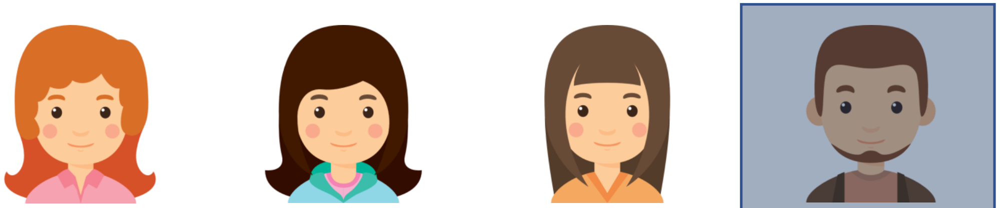

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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

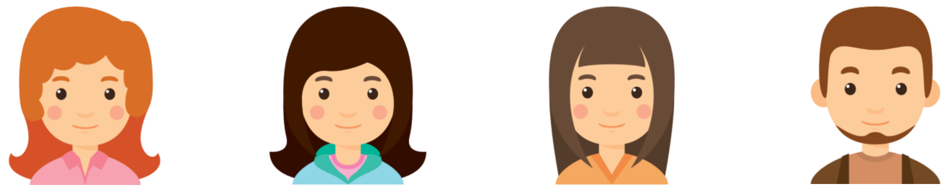

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

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 $O(m^{\frac{1}{2}-\epsilon})$ 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.
  - $v_1 \geq v_2 \geq \cdots \geq v_n$? m-approximation
- But we don't want to exhaust too many items.
  - $\frac{v_1}{|S_1|} \geq \frac{v_2}{|S_2|} \geq \cdots \frac{v_n}{|S_n|}$? 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

- OPT = Set of agents satisfied by optimal alg
- W = Set of agents satisfied by greedy alg
- For $i \in W$, let
$$OPT_i = \{j \in OPT, j \geq i : S_i \cap S_j \neq \emptyset\}$$
- $OPT \subseteq \bigcup_{i \in W} OPT_i$, so it suffices to show
$$\sqrt{m} \cdot v_i \geq \Sigma_{j \in OPT_i} v_j$$
- For each $j \in OPT_i : v_j \leq v_i \cdot \frac{\sqrt{|S_j|}}{\sqrt{|S_i|}}$

Proof of Approximation

- Summing over all $j \in OPT_i$:

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

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

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

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

Strategyproofness

- Agent i pays $p_i = v_{j^*} \cdot \sqrt{\frac{|S_i|}{|S_{j^*}|}}$
  - j^* is the smallest index $j > i$ such that $S_j \cap S_i \neq \emptyset$ and $S_j \cap S_k = \emptyset$ for all $k < j, k \neq i$
  - This is not an arbitrary value.
    - It is the lowest $\tilde{v}_i$ that agent i can report, and still win.
    - With a lower value, j^* goes first, wins, prevents i from winning.
    - "Critical payment"
  - Greedy rule is also monotonic: If agent i wins reporting (v_i, S_i), she also wins reporting $v_i' > v_i$ and $S_i' \subset S_i$.
- Critical payment + monotonic $\Rightarrow$ SP

Take-Away

- VCG can sometimes be too difficult to implement
  - Find a monotonic allocation rule that approximately maximizes welfare
  - Charge critical payments to agents

- In this case, we used approximation for computational reasons
  - In facility location, we used approximation because we couldn't use monetary payments to get SP