CS 598:
AI Methods for Market Design

Lecture 6: Mechanism Design II

Xintong Wang

Spring 2024

Announcements

- One paper presentation today! Three for next week
- Evaluation for paper presentation
- HW1 is out! Please start early

Outline

- Recap: design desiderata
- The VCG mechanism
- Optimal auctions

Outline

- Recap: design desiderata
- The VCG mechanism
- Optimal auctions

Mechanism Desiderata

- Pareto optimality
- Allocative efficiency
- Strategy proofness
- Individual rationality
- No deficit
- Budget balance

equilibrium

equilibrium strategy

Outline

- Recap: design desiderata
- The VCG mechanism
- Optimal auctions

VCG Mechanism

- The Vickrey-Clarke-Groves (VCG) mechanism
- A DRM that achieves many good properties
  - Strategy-proof (incentive compatible)
  - Allocative efficient (welfare maximizing)
  - Individually rational

VCG Mechanism

Given reported valuation profile $\hat{v} = (\widehat{v_1}, \dots, \widehat{v_n})$, the VCG mechanism on a set of alternatives A is defined by

- A choice rule

$$x(\hat{v}) = \mathrm{argmax}_{a \in A} \sum_{i \in N} \widehat{v_i}(a)$$

  with selected alternative $\mathrm{a}^* = x(\hat{v})$

VCG Mechanism

Given reported valuation profile $\hat{v} = (\widehat{v_1}, \dots, \widehat{v_n})$, the VCG mechanism on a set of alternatives A is defined by

- A choice rule

$$x(\hat{v}) = \text{argmax}_{a \in A} \sum_{i \in N} \widehat{v_i}(a)$$

  with selected alternative $\text{a}^* = x(\hat{v})$

- A payment rule: charge agent i

$$t_i(\hat{v}) = \max_{\text{a}^{-i} \in A^{-i}} \sum_{j \neq \text{i}} \widehat{v_j}\left(a^{-i}\right) - \sum_{j \neq \text{i}} \widehat{v_j}(a^*)$$

Opportunity cost incurred by agent i = The max total value to others without agent i − The total value to others under $a*$ without agent i

A^{-i} denotes the set of alternatives when agent i is not present

VCG Mechanism

Example: VCG mechanism on a single item

- Alternatives: “do not allocation” & “allocate to each agent”
- Three agents with their bids \$10, \$8, \$4 for the item
- The choice rule is $x(\hat{v}) = \text{argmax}_{a \in A}\left(\widehat{v_1}(a) + \widehat{v_2}(a) + \widehat{v_3}(a)\right)$
- The payment rule

Agent 1: $t_1(\hat{v})$ = max total value w/o 1 – current total value w/o 1

= 8 - 0 = 8

Agent 2: $t_2(\hat{v})$ = max total value w/o 2 – current total value w/o 2

= 10 - 10 = 0

VCG Mechanism

Example: VCG mechanism on a single item

Second-price auction

- Alternatives: “do not allocation” & “allocate to each agent”
- Three agents with their bids \$10, \$8, \$4 for the item
- The choice rule is $x(\hat{v}) = \text{argmax}_{a \in A}\left(\widehat{v_1}(a) + \widehat{v_2}(a) + \widehat{v_3}(a)\right)$
- The payment rule

Agent 1: $t_1(\hat{v})$ = max total value w/o 1 – current total value w/o 1

= 8 - 0 = 8

Pivotal: $a^{-i} \neq a^*$

Agent 2: $t_2(\hat{v})$ = max total value w/o 2 – current total value w/o 2

= 10 - 10 = 0

Non-pivotal: $a^{-i} = a^*$

VCG Mechanism

Example: VCG mechanism, scheduling

	9am	10am	11am
Agent 1	-5	1	2
Agent 2	20	5	10
Agent 3	5	11	2

What would be the selected alternative?

What would be the payment for each agent?

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

Proof (strategy-proof): Being truthful is dominant strategy.

1) Fix other reports v_{-i}; a is the selected alternative under (v_i, v_{-i})

$$v_i(a) - \left(\max_{\mathrm{a}^{-\mathrm{i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right) - \sum_{j \neq i} v_j(a) \right) = \sum_i v_i(a) - \max_{\mathrm{a}^{-\mathrm{i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right)$$

2) Fix other reports v_{-i}; a' is the selected alternative under $({v_i}', v_{-i})$

$$v_i(a') - \left(\max_{\mathrm{a}^{-\mathrm{i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right) - \sum_{j \neq i} v_j(a') \right) = \sum_i v_i(a') - \max_{\mathrm{a}^{-\mathrm{i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right)$$

$$\sum_i v_i(a) - \sum_i v_i(a') = \max_{\mathrm{a} \in A} \left(v_i(a) + \sum_{j \neq \mathrm{i}} v_j(a) \right) - \left(v_i(a') + \sum_{j \neq \mathrm{i}} v_j(a') \right) \geq 0$$

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

Proof (allocative efficiency):

This is by construction: $x(\hat{v}) = \mathrm{argmax}_{a \in A} \sum_i v_i(a)$

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

Proof (individual rationality): Agent i's utility of truthfulness

$$v_i(a) - \left(\max_{\mathrm{a}^{-\mathrm{i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right) - \sum_{j \neq i} v_j(a) \right)$$

$$= \sum_i v_i(a) - \max_{\mathrm{a}^{-\mathrm{i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right)$$

$$\geq \sum_i v_i(a) - \max_{\mathrm{a}^{-\mathrm{i}} \in A^{-i}} \sum_i v_i\left(a^{-i}\right)$$

$$\geq 0$$

VCG Mechanism

Recap with some more comments:

- Choose a welfare maximizing outcome
- Charge each agent i the welfare had agent i not participate *minus* the welfare of everyone else given agent i participates
- Charge each agent the "harm" it does on the welfare of everyone else (i.e., *externality*)

Computational Aspects of VCG

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay $5
- B wants one apple and is willing to pay $2
- C wants two apples and is willing to pay $6 for both but is uninterested in buying one without the other

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay $5
- B wants one apple and is willing to pay $2
- C wants two apples and is willing to pay $6 for both but is uninterested in buying one without the other

Outcome:

A and B get the two apples

A pays $4 = $6 (max value w/o A) - $2 (current value w/o A)

B pays $1 = $6 (max value w/o B) - $5 (current value w/o B)

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay \$5
- B wants one apple and is willing to pay \$2
- C wants two apples and is willing to pay \$6 for both but is uninterested in buying one without the other

What is the computational complexity of finding the welfare-maximizing outcome?

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay $5
- B wants one apple and is willing to pay $2
- C wants two apples and is willing to pay $6 for both but is uninterested in buying one without the other

What is the computational complexity of finding the welfare-maximizing outcome? A knapsack problem!

$$\text{maximize} \sum_{i=1}^{n} v_i x_i$$

$$\text{subject to} \sum_{i=1}^{n} w_i x_i \leq W \text{ and } x_i \in \{0,1\}$$

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay $5
- B wants one apple and is willing to pay $2
- C wants two apples and is willing to pay $6 for both but is uninterested in buying one without the other

What is the computational complexity of finding the welfare-maximizing outcome? A knapsack problem!

$$\text{maximize} \sum_{i=1}^{n} v_i x_i$$

$$\text{subject to} \sum_{i=1}^{n} w_i x_i \leq W \text{ and } x_i \in \{0, 1\}$$

$n = 3$ bidders
$v_1 = 5, v_2 = 2, v_3 = 6$
$w_1 = 1\ w_2 = 1, w_3 = 2$
$W = 2$

Computational Aspects of VCG

- Compute the welfare-maximizing outcome is NP-hard

- Communication cost of each agent's valuation function
  - Allocate m items to n participants
  - Each agent's valuation function consists 2^m numbers
  - Communication requirement is then $n2^m$

Outline

- Recap: design desiderata
- The VCG mechanism
- Optimal auctions

Optimal Auctions

- Efficient auctions so far
- What about maximizing the seller's expected revenue?
    - May be willing to risk failing to sell the item

Optimal Auctions: Setting

Auctions in an independent private value settings

- Risk-neutral bidders with private valuations
- Each agent's valuation v_i is independently drawn from a strictly increasing cdf $F_i(\cdot)$ with a continuous pdf $f_i(\cdot)$
  - Allow $F_i \neq F_j$: asymmetric auctions
- The risk-neutral seller knows each F_i and has no value for the object

Optimal Auctions: Definition

The auction that maximizes the expected revenue *subject to* individual rationality and Bayesian incentive compatibility for the buyers is an optimal auction

Primer: One Bidder and One Item

- A posted-price mechanism: The seller announces a price r, and the buyer can either pay the price and take the item or pay nothing and get nothing

- Given a private valuation v, how should we set a *welfare-maximizing* posted price r?

Primer: One Bidder and One Item

- A posted-price mechanism: The seller announces a price r, and the buyer can either pay the price and take the item or pay nothing and get nothing

- Given a private valuation v, how should we set a *welfare-maximizing* posted price r? $r = 0$

Primer: One Bidder and One Item

- A posted-price mechanism: The seller announces a price r, and the buyer can either pay the price and take the item or pay nothing and get nothing

- Given a private valuation v, how should we set a *welfare-maximizing* posted price r? $r = 0$

- Given a private valuation v, how should we set a *revenue-maximizing* posted price r?

Primer: One Bidder and One Item

- Given a private valuation v, how should we set a *revenue-maximizing* posted price r?

  We assume $v \sim F$, where the distribution is known but the realization v is private

  The expected revenue: $r \cdot (1 - F(r))$

  Revenue of a sale

  Prob of a sale = Pr($v \geq r$)

  When $\mathrm{v} \sim U[0,1], F(x) = x. r = \frac{1}{2}$ and E[rev]= $\frac{1}{4}$

  The **monopoly price** of F

Primer: Two Bidders in a SPA

- Given $v_i \sim U[0,1]$, what is the revenue from a welfare-maximizing auction (e.g., SPA)?

  $E[\text{rev}] = E\left[v_{(2)}\right] = \frac{1}{3}$

- Can we do better by setting a reserve price?

Primer: Two Bidders in a SPA

- How to find the optimal reserve price r^*?
  - When both $v_i < \mathrm{r}$, no sale and the revenue is 0
  - When exactly one $v_i \geq \mathrm{r}$, the revenue is r
  - When both $v_i > \mathrm{r}$, sale at second highest bid
- The dominant strategy is still to bid true value

Primer: Two Bidders in a SPA

- How to find the optimal reserve price r^*?
  - When both $v_i < r$, no sale and the revenue is 0

  Get revenue = 0 with probability r^2
  - When exactly one $v_i \geq \mathrm{r}$, the revenue is r

  Get revenue = r with probability $2(1-r)r$
  - When both $v_i > \mathrm{r}$, sale at second highest bid

  Get revenue = $E[\min v_i \mid v_i \geq r]$ with probability $(1-r)^2$
- The dominant strategy is still to bid true value

$$E[\text{rev}] = 0 \cdot r^2 + r \cdot 2(1-r)r + \frac{1+2r}{3} \cdot (1-r)^2$$

$r^* = \frac{1}{2}$ (the same reserve as one bidder one item)

Primer: Two Bidders in a SPA

- Given $v_i \sim U[0,1]$, what is the revenue from a welfare-maximizing auction (e.g., SPA)?

$$E[\text{rev}] = E\left[v_{(2)}\right] = \frac{1}{3}$$

- Can we do better by setting a reserve price?

$$E[\text{rev}] = \frac{1}{4} \cdot 0 + \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{4} \cdot \frac{2}{3} = \frac{5}{12}$$

  - Tradeoffs: higher revenue but also a risk of no sale
  - Like adding another bidder to increase competition

Primer: Two Bidders in a SPA

- Given $v_i \sim U[0,1]$, what is the revenue from a welfare-maximizing auction (e.g., SPA)?

  $$E[\text{rev}] = E\left[v_{(2)}\right] = \frac{1}{3}$$

- Can we do better by setting a reserve price?

  $$E[\text{rev}] = \frac{1}{4} \cdot 0 + \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{4} \cdot \frac{2}{3} = \frac{5}{12}$$

  - Tradeoffs: higher revenue but also a risk of no sale
  - Like adding another bidder to increase competition
- Can we do better with a different auction?

Designing Optimal Auctions

If an agent i has valuation $v_i \sim F_i$, then the *virtual valuation function* is

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

where $F(\cdot)$ and $f(\cdot)$ *are the cdf and pdf*

Designing Optimal Auctions

If an agent i has valuation $v_i \sim F_i$, then the *virtual valuation function* is

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

where $F(\cdot)$ and $f(\cdot)$ *are the cdf and pdf*

Ideally what you'd charge agent i

"information rent": revenue loss caused by not knowing v_i
Pr($v \geq v_i$) / Pr($v = v_i$)

Designing Optimal Auctions

If an agent i has valuation $v_i \sim F_i$, then the *virtual valuation function* is

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

where $F(\cdot)$ and $f(\cdot)$ *are the cdf and pdf*

1) A mapping from value space to another

- $\phi_i(v_i) \leq v_i$ and can be negative
- E.g., F is U[0, 1], then $\phi_i(v_i) = v_i - \frac{1-v_i}{1} = 2v_i - 1$

Designing Optimal Auctions

If an agent i has valuation $v_i \sim F_i$, then the *virtual valuation function* is

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

where $F(\cdot)$ and $f(\cdot)$ *are the cdf and pdf*

2) We focus on $\phi(v)$ that is monotone nondecreasing in v for all v

- Examples of regular distributions: uniform, exponential, lognormal…

Designing Optimal Auctions

If an agent i has valuation $v_i \sim F_i$, then the *virtual valuation function* is

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

where $F(\cdot)$ and $f(\cdot)$ *are the cdf and pdf*

3) We define the bidder i's bidder-specific reserve price r_i^* as the value for which $\phi_i(r_i^*) = 0$

- E.g., F is U[0, 1], then $\phi_i(v_i) = 2v_i - 1$ and $r_i^* = \frac{1}{2}$

Designing Optimal Auctions

Lemma. In *single-dimensional* settings, where there are n agents with private values v_i drawn independently from *known* distributions F_i.

For every strategy-proof, DRM $M = (x, t)$, for every agent i:

$$E_{v_i \sim F_i}[t_i(v)] = E_{v_i \sim F_i}[\phi_i(v_i)x_i(v)]$$

Expected payment for agent i, given input v	=	Expected virtual value for agent i, given input v

Designing Optimal Auctions

Lemma. […] For every strategy-proof, DRM $M = (x, t)$, for every agent i:

$$E_{v_i \sim F_i}[t_i(v)] = E_{v_i \sim F_i}[\phi_i(v_i) x_i(v)]$$

Proof (sketch): utilize the payment identity

Recap: Characterizing BNE in Auctions

Theorem. In any BNE of any sealed-bid auction, for bidder i with v_i, we have

- Interim monotonicity: The interim allocation $x_i^*(v_i)$ is monotone weakly increasing in value v_i
- Interim payment identity: For value v_i and interim allocation $x_i^*(v_i)$, the interim payment is

$$t_i^*(v_i) = v_i \times x_i^*(v_i) - \int_{z=v_{\min}}^{v_i} x_i^*(z)\mathrm{d}z - C_i$$

  where C_i is a constant.

Rate of increase in interim payment = Rate of increase in interim value

Designing Optimal Auctions

Lemma. [...] For every strategy-proof, DRM $M = (x, t)$, for every agent i:

$$E_{v_i \sim F_i}[t_i(v)] = E_{v_i \sim F_i}[\phi_i(v_i) x_i(v)]$$

Proof (sketch): utilize the payment identity

$$\begin{aligned}
\mathbb{E}_{v_i \sim F_i}(t_i(v_i, v_{-i})) &= \int_0^{v_{\max}} t_i(z, v_{-i}) f_i(z) dz \\
&= \int_0^{v_{\max}} \int_0^{z} \underline{y \cdot x_i'(y, v_{-i})} f_i(z) dy dz \\
&= \int_0^{v_{\max}} \left(\frac{1 - F_i(y)}{f_i(y)} \right) f_i(y) x_i(y, v_{-i}) dy \\
&= \mathbb{E}_{v_i \sim F_i}(\phi(v_i) x_i(v)).
\end{aligned}$$

Designing Optimal Auctions

Theorem. In *single-dimensional* settings, for every strategy-proof, DRM $M = (x, t)$:

$$E_{v \sim F}\left[\sum_i t_i(v)\right] = E_{v \sim F}\left[\sum_i \phi_i(v_i) x_i(v)\right]$$

Expected revenue of M, given input v	=	Expected total virtual value of allocation, given input v

- Maximizing expected revenue is equivalent to maximizing the virtual welfare!
- We want to find the virtual welfare maximizing $x^*(v)$

Designing Optimal Auctions

Intuition:

- If all virtual values are negative, do not allocate
- Otherwise, allocate the item to the bidder with the highest virtual value $\phi_i(v_i)$ (possibly not highest value bidder!)
- By how much should we charge the winning bidder?

Myerson’s Optimal Auctions

Theorem (Myerson 1981) Suppose there are n bidders with valuations $v_i \sim F_i$ drawn independently from regular distributions and a single item for sale. The revenue-optimal, incentive compatible auction in terms of a DRM:

- Allocate the item to agent $i = \text{argmax}_i \phi_i(\hat{v_i})$ if $v_i \geq r_i^*$
- If a sale, charge the winning agent i the smallest valuation that it could have declared while remaining the winner, i.e.,

$$\inf\{v_i^*\colon \phi_i(v_i^*) \geq 0 \text{ and } \phi_i(v_i^*) \geq \phi_j(\hat{v_j}), \forall i \neq j\}$$

Myerson's Optimal Auctions

Corollary (Myerson 1981) In a symmetric setting where there are n bidders, the optimal, incentive compatible auction is a SPA with a reserve price of r^* that solves $r^* - \frac{1-F(r^*)}{f(r^*)} = 0$

Myerson's Optimal Auctions

Corollary (Myerson 1981) In a symmetric setting where there are n bidders, the optimal, incentive compatible auction is a SPA with a reserve price of r^* that solves $r^* - \frac{1-F(r^*)}{f(r^*)} = 0$

Let's verify with prior examples

- One bidder one item: $r \cdot \left(1 - F(r)\right)$

$1 - F(r) - rf(r) = 0$ → $r - \frac{1-F(r)}{f(r)} = 0$ → $r = \phi^{-1}(0)$

Analyzing Optimal Auctions

Optimal Auction:

winning agent $i = \text{argmax}_i \phi_i(\hat{v}_i)$ if $v_i \geq r_i^*$

Agent i is charged the smallest valuation that it could have declared while remaining the winner, i.e.,

$$\inf\{v_i^*:\ \phi_i(v_i^*) \geq 0 \text{ and } \phi_i(v_i^*) \geq \phi_j\left(\hat{v}_j\right), \forall i \neq j\}$$

- Is this VCG? *No, it's not efficient.*
- How should bidders bid?
  - It's a SPA with a reserve price, held in virtual valuation space
  - Neither r_i^* nor $\phi_i(v_i)$ depends on the agent report
  - Thus, the proof that a SPA is strategy-proof holds