TTIC 31290: Machine Learning for Algorithm Design (Fall 2025)
Avrim Blum and Dravyansh Sharma

Lecture 10: 11/04/25 Lecturer: Dravyansh Sharma

Outline for today

- Second-price Auctions
- Multi-item generalization (VCG Mechanism)
- Data-driven mechanism design

1 Second-price Auctions

We consider a single indivisible item for sale among n bidders. Each bidder $i \in \{1, 2, \ldots, n\}$ has a private valuation $v_i \geq 0$ for the item, representing the maximum amount they are willing to pay.

- Each bidder submits a sealed bid $b_i \geq 0$.
- Let $b_{(1)} = \max_i b_i$ denote the highest bid, and $b_{(2)}$ the second-highest bid.

A **second-price (Vickrey) auction** is defined by the following allocation and payment rules:

(a) The bidder i^* with the highest bid wins the item:

$$i^* = \arg\max_i b_i.$$

(b) The winner pays the *second-highest* bid as the price:

$$p_{i^*} = b_{(2)}.$$

(c) All other bidders $j \neq i^*$ receive nothing and pay zero:

$$p_j = 0.$$

For a bidder i, define the **utility** function under quasi-linear preferences as:

$$u_i(b_i, b_{-i}) = \begin{cases} v_i - p_i & \text{if bidder } i \text{ wins the item,} \\ 0 & \text{otherwise.} \end{cases}$$

1.1 Second-price auctions are strategy-proof

Definition 1. *A bidding strategy $b_i(v_i)$ is a* ***dominant strategy*** *for bidder i if for every possible valuation profile of the other bidders v_{-i} and their corresponding bids b_{-i},*

$$u_i(b_i(v_i), b_{-i}) \geq u_i(b'_i, b_{-i}) \quad \textit{for all possible deviations } b'_i.$$

We will now show that truthful bidding maximizes bidder i's utility regardless of what others do.

Theorem 1. *In a second-price auction, it is a dominant strategy for each bidder i to bid their true valuation:*

$$b_i = v_i.$$

Proof. Fix any bidder i with valuation v_i, and let b_{-i} be the bids of all other bidders. Let $b_{(1)}^{-i} = \max_{j \neq i} b_j$ denote the highest bid among others.

(a) **Case 1:** $v_i < b_{(1)}^{-i}$.

If i bids truthfully, $b_i = v_i < b_{(1)}^{-i}$, so i loses and utility is 0. If i overbids (i.e., $b_i > b_{(1)}^{-i}$), i wins and pays $b_{(1)}^{-i}$, obtaining

$$u_i = v_i - b_{(1)}^{-i} < 0.$$

Thus overbidding strictly reduces utility. Underbidding changes nothing. Hence truthful bidding maximizes utility.

(b) **Case 2:** $v_i > b_{(1)}^{-i}$.

If i bids truthfully, $b_i = v_i > b_{(1)}^{-i}$, so i wins and pays $b_{(1)}^{-i}$, getting

$$u_i = v_i - b_{(1)}^{-i} > 0.$$

If i underbids (i.e., $b_i < b_{(1)}^{-i}$), i loses and utility is 0, which is worse. Overbidding does not change the outcome, since i would still win and pay $b_{(1)}^{-i}$.

Therefore, bidding $b_i = v_i$ weakly dominates all other strategies. □

Remark 1. *The second-price auction is an example of a* dominant-strategy incentive compatible *(DSIC) mechanism by Theorem 1. It is also* efficient*, the item is always allocated to the bidder with the highest valuation.*

2 Multi-item generalization (VCG Mechanism)

We now consider the generalization of the second-price auction to the case of multiple items or, more generally, arbitrary outcomes.

Let:

- $\mathcal{O}$ be the set of possible *outcomes* (e.g., allocations of items).

- There are n bidders, indexed by $i = 1, \ldots, n$.
- Each bidder i has a private **valuation function**

$$v_i : \mathcal{O} \to \mathbb{R}_{\geq 0},$$

  representing how much i values each possible outcome.

An example special case is a **multi-item auction**, where $\mathcal{O}$ is the set of allocations of m items among n bidders, and $v_i(S_i)$ is the value bidder i assigns to receiving the subset S_i of items.

2.1 Mechanism Definition

Each bidder reports a bid function $b_i : \mathcal{O} \to \mathbb{R}_{\geq 0}$, which may differ from their true valuation v_i.

Outcome Rule (Allocation). The mechanism selects the outcome that maximizes total reported welfare:

$$o^*(b) = \arg\max_{o \in \mathcal{O}} \sum_{i=1}^{n} b_i(o).$$

Payment Rule. Each bidder i pays

$$p_i(b) = \left[\max_{o \in \mathcal{O}} \sum_{j \neq i} b_j(o) \right] - \sum_{j \neq i} b_j(o^*(b)).$$

That is, bidder i pays the *externality* they impose on others,

$$p_i = (\text{welfare of others without } i) - (\text{welfare of others with } i).$$

Bidder i's utility is

$$u_i(b) = v_i(o^*(b)) - p_i(b).$$

2.2 VCG is DSIC

Theorem 2. *In the Vickrey–Clarke–Groves mechanism, truthful reporting is a dominant strategy for each bidder,*

$$b_i = v_i.$$

Proof. Fix any bidder i and any reports b_{-i} of others.

Let $o^*(v_i, b_{-i})$ denote the outcome chosen if i bids truthfully, and $o^*(b'_i, b_{-i})$ denote the outcome if i misreports as b'_i.

Under the VCG payment rule

$$u_i(b'_i, b_{-i}) = v_i(o^*(b'_i, b_{-i})) - \left(\left[\max_{o \in \mathcal{O}} \sum_{j \neq i} b_j(o) \right] - \sum_{j \neq i} b_j(o^*(b'_i, b_{-i})) \right).$$

The term $\left[\max_{o\in\mathcal{O}} \sum_{j\neq i} b_j(o)\right]$ does not depend on b'_i and thus does not affect the optimization. Therefore, bidder i maximizes

$$v_i(o^*(b'_i, b_{-i})) + \sum_{j\neq i} b_j(o^*(b'_i, b_{-i})).$$

But $o^*(b'_i, b_{-i})$ is defined as the maximizer of $\sum_k b_k(o)$. When $b'_i = v_i$, this expression becomes exactly

$$o^*(v_i, b_{-i}) = \arg\max_o \sum_k v_k(o),$$

i.e., the outcome maximizing true total welfare.

Thus, by bidding truthfully, bidder i ensures that the chosen outcome maximizes their own utility function above. Any misreport can only reduce this quantity, so $b_i = v_i$ is a dominant strategy. □

We note the following.

- When there is a *single item*, the VCG mechanism reduces to the **second-price auction**.
- When there are multiple items but additive valuations ($v_i(S_i) = \sum_{j\in S_i} v_{ij}$), the VCG mechanism corresponds to each item being sold via an independent second-price auction.
- For combinatorial valuations, VCG ensures efficiency and truthfulness but may be computationally intractable (the outcome rule can be NP-hard to compute).

3 Data-driven mechanism design

Suppose there are m distinct items, and the bidders have additive valuations for the different items. We consider a slight extension of the above to include *reserve prices*. For each item j, we set a minimum price p_j so that the item is sold for at least that price. The bidder with the highest bid for an item j must have bid at least p_j to get the item (else no one gets item j), and their payment is the maximum of the second highest bid and the reserve price p_j. For fixed reserve prices, the mechanism is still DSIC.

Suppose there is an unknown distribution D over the bidders' values. Since VCG is DSIC, so we assume that the bids equal the bidders' valuations. Can we design a mechanism that sets the reserve prices to maximize the expected revenue (total payment for all items) over D, given access to sample auctions?

Theorem 3. *The revenue function class, consisting of functions parameterized by the reserve prices* $\mathbf{p} = (p_1, \ldots, p_m)$ *that give the revenue of any auction for a given reserve prices, is* $(\mathcal{F}, \mathcal{G}, 2m)$*-decomposable, where* $\mathcal{F}$ *consists of linear functions* $\mathbb{R}^m \to \mathbb{R}$ *and* $\mathcal{G}$ *consists of linear thresholds* $\mathbb{R}^m \to \{0, 1\}$*.*

Proof. Given a fixed instance (fixed valuation and fixed bids), let i_j and i'_j be the highest and second highest bidders for item j. Item j is sold to bidder i_j if the bid $b_{i_j} \geq p_j$ and the revenue is either p_j or $b_{i'_j}$ depending on whether $p_j \geq b_{i'_j}$. Thus, there are $2m$ axis-aligned hyperplanes that partition $\mathbb{R}^m$ into regions where the revenue is linear. □

Additional Resources:

- Tim Roughgarden. Twenty lectures on algorithmic game theory. Cambridge University Press, 2016.
- Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. "Generalization guarantees for multi-item profit maximization: Pricing, auctions, and randomized mechanisms." Operations Research 73, no. 2 (2025): 648-663.