DS 574 Algorithmic Mechanism Design	**Lecture #6**
Fall 2023	Prof. Kira Goldner

Welfare Maximization in Multidimensional Settings

Multidimensional or *multi-parameter* environments are ones where we need to elicit more than one piece of information per bidder. The most common settings include m heterogenous (*different*) items and

- n unit-demand buyers; buyer i has value v_{ij} for item j but only wants at most 1 item. (You only want to buy 1 house!)
- n additive buyers: buyer i's value for set S is $\sum_{j \in S} v_{ij}$.
- n subadditive buyers for some subadditive functions
- n buyers who are k-demand: buyer i's value for a set of items S is $\max_{|S'|=k, S' \subseteq S} \sum_{j \in S'} v_{ij}$.
- n matroid-demand buyers for some matroid
- ...

With m heterogenous items, it's *possible* that our buyers could have different valuations for every single one of the 2^m bundles of items—that is why this general setting is referred to as *combinatorial auctions.*

Then how can we maximize welfare in this setting? How can we do so *tractably?* How can we even elicit preferences in a tractable way?

Theorem 1 (The Vickrey-Clarke-Groves (VCG) Mechanism)**.** *In every general mechanism design environment, there is a DSIC welfare-maximizing mechanism.*

Given bids $\mathbf{b}_1, \ldots, \mathbf{b}_n$ where each bid is indexed by the possible outcomes Ω, we define the welfare-maximizing allocation rule $\mathbf{x}$ by

$$\mathbf{x}(\mathbf{b}) = \text{argmax}_{\omega \in \Omega} \sum_{i=1}^{n} b_i(\omega).$$

Now that things are multidimensional, there's no more Myerson's Lemma! In multiple dimensions, what is monotonicity? What would the critical bid be?

Instead, we have bidders pay their *externality*—the loss of welfare caused due to i's participation:

$$p_i(\mathbf{b}) = \underbrace{\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega)}_{\text{without } i} - \underbrace{\sum_{j \neq i} b_j(\omega^*)}_{\text{with } i}$$

where $\omega^* = \mathbf{x}(\mathbf{b})$ is the outcome chosen when i *does* participate.

Claim 1. The VCG mechanism is DSIC.

Proof. We show that the mechanism with (x, p) is DSIC: that setting $\mathbf{b}_i = \mathbf{v}_i$ maximizes utility $v_i(\mathbf{x}(\mathbf{b})) - p_i(\mathbf{b})$. Fix i and $\mathbf{b}_{-i}$.

When the chosen outcome $\mathbf{x}(\mathbf{b})$ is ω^*, i's utility is

$$v_i(\omega^*) - p_i(\mathbf{b}) = \left[v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*)\right] - \left[\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega)\right].$$

The second term is independent of i's bid. The first term is equal to social welfare, which x is chosen to maximize for the input bids. Thus the mechanism is aligned with i's incentives, and i's utility is maximized when i reports their true valuations. □

Exercise (optional): Prove that the payment $p_i(\mathbf{b})$ is always non-negative (and so the mechanism is IR).

Proof. The outcome in the first term of the payment is chosen to maximize it, whereas the second term is the same but not with the optimal outcome for the term, hence the first term is larger. □

Ascending Auctions

In *ascending auctions*, an auctioneer initializes prices for each item, iteratively raises the prices, and bidders decide which items to bid on in each round. Sometimes *activity rules* are enforced, e.g., once you drop out on an item, you can not bid on it again.

The most famous ascending auction is the single-item version, the English Auction.

The English Auction(ε):

a. Initialize the item's price p_0 to 0.

b. The initial set S_0 of "active bidders" (willing to pay p_0 for the item) is all bidders.

c. For iteration $t = 1, 2, \ldots,$:

   (a) Ask the set of active bidders S_{t-1} if they're willing to pay $p_{t-1} + \varepsilon$. Let S_t be the bidders who say yes. (Hopefully, $v_i \geq p_{t-1} + \varepsilon$.)

   (b) If $|S_t| \leq 1$: terminate the auction. Allocate the item to the remaining active bidder at a price of p_{t-1}. If no bidders remain, randomly allocate to a bidder from S_{t-1} at p_{t-1}.

   (c) Otherwise, $p_t = p_{t-1} + \varepsilon$.

Benefits of using ascending auctions:

- Ascending auctions are easier for bidders. It is generally easier to answer simple queries than to report a valuation. This point will become especially relevant in more complex scenarios.
- Less information leakage. The winner of an ascending auction does not reveal its valuation, just the fact that it is at least the second-highest bid.

- Transparency. The cause of a high selling price is generally more obvious in open ascending auctions than in sealed-bid auctions.
- Potentially more seller revenue. For example, ascending auctions encourage "bidding wars." There is also some supporting theoretical work on this point [1].
- When there are multiple items, the opportunity for "price discovery." A bidder has the opportunity for mid-course corrections and to better coordinate with other bidders.

What about k identical items? What should we do here?

The English Auction for k Identical Items:

The same as above, but replace step 3(b) with the following:

(b) If $|S_t| \leq k$: terminate the auction. Allocate the items to the remaining active bidders at a price of p_{t-1}. If there are items leftover (i.e., $k - |S_t| > 0$), randomly allocate them to bidders from $S_{t-1} \setminus S_t$ at p_{t-1}.

Definition 1. In an ascending auction, *sincere bidding* means that a player answers all queries honestly.

Claim 2. In the k identical item setting, in an English auction, sincere bidding is a dominant strategy for every bidder (up to ε).

Claim 3. In the k identical item setting, if all bidders bid sincerely in an English auction, the welfare of the outcome is within $k\varepsilon$ of the maximum possible.

The English auction for k Identical Items terminates in $v_{\max}/\varepsilon$ iterations.

The above claims are left as an exercise.

We can use the following design process for ascending auctions:

a. As a sanity check, design a direct-revelation DSIC welfare-maximizing polytime mechanism.

b. Implement this as an ascending auction.

c. **(Truthfulness)** Check that its EPIC.

d. **(Performance)** Check that it still maximizes welfare under sincere bidding.

e. **(Tractability)** Check that it terminates in a reasonable number of iterations.

Additive Valuations, Parallel Auctions

The Additive Setting: There are m non-identical items and n bidders where each bidder i has private valuation v_{ij} for each item j. Bidder i has an additive valuation for each set S, that is,

$$v_i(S) := \sum_{j \in S} v_{ij}.$$

Step 1: What is the welfare-optimal direct revelation mechanism here? Just handle each item separately—m Vickrey auctions!

What's the analogous ascending implementation?

Parallel English Auctions: Maintain a set of interested bidders for each item, and the auction for item j terminates when there's only one active bidder remaining, breaking ties arbitrarily.

Is this DSIC? No!

Example: Two bidders, two items. $\mathbf{v}_1 = (3, 2)$ and $\mathbf{v}_2 = (2, 1)$.

What happens under sincere bidding? The first bidder wins both items at prices of 2 and 1 respectively.

Alternatively, bidder 2 could threaten the following strategy: if bidder 1 bids on item 1 in the first turn, then bidder 2 will keep bidding on both items forever (or up to a price of 3). If not, they will bid sincerely until the auction terminates.

Then bidder 1 bidding sincerely triggers bidder 2's threat, causing bidder 1 to lose both items, so bidder 1 would prefer to abandon item 1.

Recall that a dominant strategy maximizes a bidder's utility independent of the actions played by any other player. Bidder 2's strategy may not maximize their utility, but it still implies that sincere bidding is not a *dominant* strategy for bidder 1.

Instead, we need a different solution concept.

Definition 2. A strategy profile $(\sigma_1, \ldots, \sigma_n)$ is an *ex post Nash equilibrium (EPNE)* if, for every bidder i and valuation $v_i \in V_i$, the strategy $\sigma_i(v_i)$ is a best-response to every strategy profile $\sigma_{-i}(\mathbf{v}_{-i})$ with $\mathbf{v}_{-i} \in \mathbf{V}_{-i}$.

In comparison, in a dominant-strategy equilibrium (DSE), for every bidder i and valuation v_i, the action $\sigma_i(v_i)$ is a best response to every action profile $\mathbf{a}_{-i}$ of $\mathbf{A}_{-i}$, whether of the form $\sigma_{-i}(\mathbf{v}_{-i})$ or not.

Definition 3. A mechanism is *ex post incentive compatible (EPIC)* if sincere bidding is an ex post Nash equilibrium in which all bidders always receive nonnegative utility.

Claim 4. For n additive bidders with m heterogenous items, in parallel English auctions, sincere bidding by all bidders is an ex post Nash equilibrium (up to $m\varepsilon$).

Acknowledgements

This lecture was developed in part using materials by Tim Roughgarden, and in particular, his book "Twenty Lectures on Algorithmic Game Theory" and lectures from "CS364B: Frontiers in Mechanism Design" [1, 2].

References

[1] Tim Roughgarden. CS364B: Frontiers in Mechanism Design, 2014.

[2] Tim Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016.