DS 574 Algorithmic Mechanism Design	**Lecture #6**
Fall 2025	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. $\square$

What does the VCG mechanism look like for:

- bidders with additive valuations? $v_i(S) = \sum_{j \in S} v_{ij}$. The VCG mechanism here is just m (number of items) separate second-price auctions.
- unit-demand bidders? $v_i(S) = \max_{j \in S} v_{ij}$. The VCG mechanism here, for its allocation, chooses the maximum weight matching. The payment for i is then [the weight of the maximum weight matching without bidder i on the left] - [the weight of the maximum weight matching without the weight of the edge from bidder i to their matching item].

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. $\square$

Interdependent Values I

Thus far, we have been discussing private independent values. That is, each bidder i has private information $\mathbf{v}_i$ regarding their value for item i.

However, in many settings, there valuations may be correlated between buyers, depend on one another's information, or even be common.

The Interdependent Values Model [2]. Each bidder has a private *signal* s_i that is a piece of information about the item, so in total the information about the item is $s_1, \ldots, s_n$, but is distributed amongst the different buyers. Each buyer has a **public** *valuation function* $v_i(s_1, \cdots, s_n)$ that dictates how the buyer aggregates the information into a value for the item.

Assumptions on $v_i(\cdot)$:

- $v_i(\cdot)$ monotone in s_j for all i, j.
- $v_i(\cdot)$ is non-negative for all $\mathbf{s}$.

Example: Common Values [6]: The average of estimates $v_i(s_1, \ldots, s_n) = \frac{1}{n}\sum_i s_i \, \forall i$, or the wallet game $v_i(s_1, \ldots, s_n) = \sum_i s_i \, \forall i$.

Optimal Social Welfare

Mechanisms. How can we maximize social welfare in this setting, optimally? What does a mechanism even look like?

- Report: A bid of a signal b_i for each bidder i, truthful when $b_i = s_i$.
- Calculate: $v_i(\mathbf{b})$ for each bidder i
- Allocate to: [This is the decision of the mechanism.]

Incentive Compatibility. What conditions are necessary for maximizing social welfare optimally to be incentive-compatible? What definition of incentive-compatible are we going for?

Give an example showing why we can't expect our mechanisms to be DSIC.

So the next best we can hope for is EPIC. In this context that means:

Definition 1. Truth-telling is said to be ex-post Nash if, for every bidder i, for every possible realization of the other bidders' signals $\mathbf{s}_{-i}$, and given that other bidders report their signals truthfully, then it is in bidder i' best interest to report her true signal.

What is the analogue of Myerson's Lemma in the interdependent setting?

Lemma 1 (Myerson Analogue [5])**.** *For every interdependent values setting,*

(a) An allocation rule $\mathbf{x}$ is implementable as EPIC and ex post IR if and only if for every $i, \mathbf{s}_{-i}$, the allocation rule x_i is monotone non-decreasing in the signal s_i.

(b) If $\mathbf{x}$ is monotone, then there is a unique payment rule such that the sealed-bid mechanism $(\mathbf{x}, \mathbf{p})$ is EPIC and ex-post IR.

(c) The payment rule in is given by:

$$\begin{aligned} p_i(\mathbf{s}) = x_i(\mathbf{s})v_i(\mathbf{s}) &- \int_{v_i(0,\mathbf{s}_{-i})}^{v_i(s_i,\mathbf{s}_{-i})} x_i(v_i^{-1}(t \mid \mathbf{s}_{-i}), \mathbf{s}_{-i})dt \\ &- \left[x_i(0,\mathbf{s}_{-i})v_i(0,\mathbf{s}_{-i}) - p_i(0,\mathbf{s}_{-i})\right]; \\ p_i(0,\mathbf{s}_{-i}) \leq x_i(0,\mathbf{s}_{-i})&v_i(0,\mathbf{s}_{-i}). \end{aligned}$$

Derivation. Fix a bidder i with public valuation function $v_i(\cdot)$. Let $s_{i\ell}$ be the ℓ^{th} possible realization of s_i in the discrete support of i's signals. Fix the signals of the other bidders $\mathbf{s}_{-i}$, and we discuss the possible values of bidder i in the context of the support of the values $\{s_{i0} = 0, s_{i1}, \ldots, s_{ik}\}$ for some high k.

For notational brevity, in the following derivation, we drop the $\mathbf{s}_{-i}$ in the input, writing just $v_i(s_i)$, $x_i(s_i)$, and $p_i(s_i)$ instead of $v_i(s_i, \mathbf{s}_{-i})$, $x_i(s_i, \mathbf{s}_{-i})$, and $p_i(s_i, \mathbf{s}_{-i})$. Then using the fact that we seek an EPIC mechanism, we deduce the following.

The bidder with signal $s_{i\ell}$ prefers truthful reporting to reporting $s_{i\ell-1}$:

$$v_i(s_{i\ell})x_i(s_{i\ell}) - p_i(s_{i\ell}) \geq v_i(s_{i\ell})x_i(s_{i\ell-1}) - p_i(s_{i\ell-1})$$

The bidder with signal $s_{i\ell-1}$ prefers truthful reporting to reporting $s_{i\ell}$:

$$v_i(s_{i\ell-1})x_i(s_{i\ell-1}) - p_i(s_{i\ell-1}) \geq v_i(s_{i\ell-1})x_i(s_{i\ell}) - p_i(s_{i\ell})$$

Thus, this gives that:

$$\begin{aligned} v_i(s_{i\ell})\left[x_i(s_{i\ell}) - x_i(s_{i\ell-1})\right] &\geq p_i(s_{i\ell}) - p_i(s_{i\ell-1}) \\ &\geq v_i(s_{i\ell-1})\left[x_i(s_{i\ell}) - x_i(s_{i\ell-1})\right]. \end{aligned}$$

Under the assumption that $s_{i0} = 0$, this gives

$$\frac{\partial}{\partial s_i} p_i(s_i, \mathbf{s}_{-i}) \geq v_i(s_i, \mathbf{s}_{-i}) \frac{\partial}{\partial s_i} x_i(s_i, \mathbf{s}_{-i})$$

and hence

$$\begin{aligned} p_i(s_i, \mathbf{s}_{-i}) &= \int_0^{s_i} \frac{\partial}{\partial z} p_i(z, \mathbf{s}_{-i})\, dz + p_i(0, \mathbf{s}_{-i}) \\ &= \int_0^{s_i} v_i(z, \mathbf{s}_{-i}) \frac{\partial}{\partial s_i} x_i(z, \mathbf{s}_{-i})\, dz + p_i(0, \mathbf{s}_{-i}) \\ &= x_i(\mathbf{s})v_i(\mathbf{s}) - x_i(0, \mathbf{s}_{-i})v_i(0, \mathbf{s}_{-i}) - \int_0^{s_i} x_i(z, \mathbf{s}_{-i}) \frac{\partial}{\partial z} v_i(z, \mathbf{s}_{-i})\, dz + p_i(0, \mathbf{s}_{-i}) \\ &= x_i(\mathbf{s})v_i(\mathbf{s}) - \int_0^{s_i} x_i(z, \mathbf{s}_{-i}) \frac{\partial}{\partial z} v_i(z, \mathbf{s}_{-i})\, dz - \left[x_i(0, \mathbf{s}_{-i})v_i(0, \mathbf{s}_{-i}) - p_i(0, \mathbf{s}_{-i})\right] \\ &= x_i(\mathbf{s})v_i(\mathbf{s}) - \int_{v_i(0, \mathbf{s}_{-i})}^{v_i(s_i, \mathbf{s}_{-i})} x_i(v_i^{-1}(t \mid \mathbf{s}_{-i}), \mathbf{s}_{-i}) dt - \left[x_i(0, \mathbf{s}_{-i})v_i(0, \mathbf{s}_{-i}) - p_i(0, \mathbf{s}_{-i})\right]. \end{aligned}$$

And we need to also ensure ex-post individual rationality for the type with signal 0:

$$p_i(0, \mathbf{s}_{-i}) \leq x_i(0, \mathbf{s}_{-i})v_i(0, \mathbf{s}_{-i}).$$

This is typically guaranteed by setting $p(0) = 0$ in the independent private value setting, but $s_i = 0$ doesn't mean that $v_i(0, \mathbf{s}_{-i}) = 0$. Guaranteeing it for the type with signal 0 ensures it for the rest of the types by the payment identity (which ensures EPIC among types).

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" [3, 4].

References

[1] Alon Eden, Michal Feldman, Amos Fiat, and Kira Goldner. Interdependent values without single-crossing. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pages 369–369, New York, NY, USA, 2018. ACM.

[2] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.

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

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

[5] Tim Roughgarden and Inbal Talgam-Cohen. Optimal and robust mechanism design with interdependent values. *ACM Trans. Econ. Comput.*, 4(3):18:1–18:34, June 2016.

[6] Robert B Wilson. Competitive bidding with asymmetric information. *Management Science*, 13(11):816–820, 1967.