DS 574 Algorithmic Mechanism Design | **Lecture #6**
Fall 2022 | Prof. Kira Goldner

Intuition and Clarifications for Quantile Space and Ironing

Define *quantiles* mapping to values such that for any value v, $q = 1 - F(v)$ represents the fraction of the distribution with value at least v, or willing to pay a price of v. Recall:

- $q = 1 - F(v)$
- $v(q) = F^{-1}(1-q)$
- Independent of F, $q \sim U[0,1]$: the probability that a value is drawn from the top .3 is .3.
- Quantile space flips: values are decreasing in increasing quantiles, and the allocation must be weakly decreasing in increasing quantiles.

Today we'll repeat the results from last time, but in value space. Remember, however, that while we can do everything we did in quantile space instead in value space, because quantile space is uniform *always* and value space is not, we'll need to normalize by the distribution in value space.

Definition 1. The *price-posting revenue curve* of a single-dimensional linear agent specified by $R(v) = v \cdot [1 - F(v)]$. Pointwise this equal to the price-posting revenue curve in quantile space: $R(v(q)) = P(q) = v(q) \cdot q$.

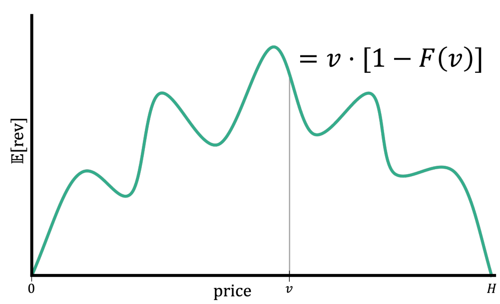

Figure 1: A price-posting revenue curve in *value space*.

Note: This is *only* the revenue that can be achieved by posting a single take-it-or-leave-it price. This does not capture the expected revenue of any given mechanism.

Claim 1. A distribution F is regular if and only if:

- its corresponding price-posting revenue curve *in quantile space* is concave.
- $\varphi(q)$ is strictly increasing.
- $f(v)\varphi(v)$ is strictly increasing.

(a) Price of \$3.

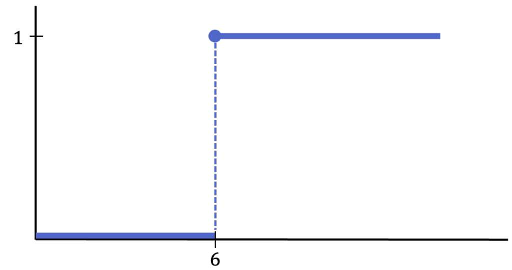

(b) Price of \$6.

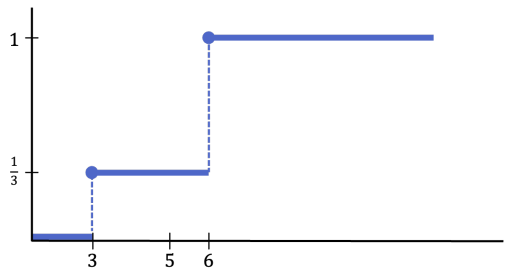

(c) Randomized price with expectation \$5.

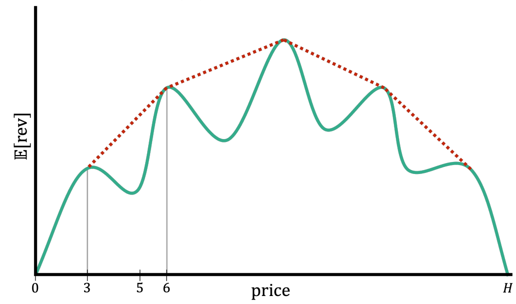

(d) Ironed revenue curve.

Figure 2: (a) An allocation rule for a take-it-or-leave-it price of \$3. (b) An allocation rule for a take-it-or-leave-it price of \$6. (c) An allocation that can be written $x(v) = 0$ for $v < 3$, $x(v) = \frac{1}{3}$ for $v \in [3, 6)$, and $x(v) = 1$ for $v \geq 6$. Alternatively, a randomized take-it-or-leave-it price that is \$3 with probability $\frac{1}{3}$ and \$6 with probability $\frac{2}{3}$, that is, $\$5 = \frac{1}{3} \cdot 3 + \frac{2}{3} \cdot 6$ in expectation. (d) The revenue curve in value space, including ironed intervals where convex combinations of prices can attain higher revenue than deterministic prices.

Claim 2. Any DSIC allocation rule $x(\cdot)$ can be expressed as a distribution of posted prices.

Recall the ironing procedure: Take the concave hull of the price-posting revenue curve in quantile space. Its derivative forms the ironed virtual values. (The derivatives of the original curve are the original virtual values.)

Claim 3. The expected revenue on the ironed revenue curve is attainable with a DSIC mechanism.

Example: How would you obtain the ironed revenue at \$5 instead of just $R(5)$?

For $p \in [\underline{p}, \overline{p}]$ where $\overline{R}(p) > R(p)$, if $p = \alpha\underline{p} + (1-\alpha)\overline{p}$, we achieve $\overline{R}(p)$ by randomizing the prices $\underline{p}$ and $\overline{p}$ with probabilities α and $1-\alpha$ accordingly to yield $\alpha R(\underline{p}) + (1-\alpha)R(\overline{p})$ on the concave closure.

Note: Recall that the expected revenue of *any mechanism*, not just a posted price, can be expressed by its virtual welfare. (We have now shown that you could decompose it into a distribution of posted prices and thus express the revenue that way, too, actually.)

Theorem 1. *For any monotone allocation rule $x(\cdot)$ and any virtual value function $\varphi(\cdot)$, the expected virtual welfare of an agent is upper-bounded by their expected ironed virtual welfare, i.e.,*

$$\mathbb{E}[\varphi(v)x(v)] \leq \mathbb{E}[\bar{\varphi}(v)x(v)].$$

Furthermore, this inequality holds with equality if the allocation rule x satisfies $x'(v) = 0$ for all v where $\bar{\Phi}(v) > \Phi(v)$.

What's the final mechanism? Now that $\bar{\varphi}_i(\cdot)$ is monotone (for every i), we choose the $x(\cdot)$ that maximizes $\mathbb{E}_v[\sum_i \varphi_i(v)x_i(v)]$, which will thus be monotone. By Theorem 1, this is an *upper bound* on the optimal revenue.

For any ironed interval $[a, b]$, examine $\bar{\varphi}(v)$ for $v \in [a, b]$. $P(q(v))$ is a straight line (linear) there, so $\bar{\varphi}(q(v))$ will be constant.

What does this imply for ironed-virtual-welfare-maximizing allocation in $[a, b]$? It will be constant on $[a, b]$, and thus its derivative will be zero.

Hence ironed virtual welfare is equal to virtual welfare by Theorem 1, so maximizing one maximizes the other.

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 2 (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 4. 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. □