Revenue Maximization

Instructor: Thomas Kesselheim

So far, our attention in mechanism design focused on social welfare. That is, we wanted to maximize the overall value of the allocation that we make. Today we move to a different objective function, namely to maximize revenue. How can we sell an item so as to maximize the winner's payment?

This question is a lot different from maximizing social welfare. For example, assume that we have only a single bidder. Maximizing social welfare is trivial (just give him the item) but how do we make him pay as much as possible? If we have no idea of what the item could be worth to him, he can just claim arbitrarily small numbers. Therefore, the standard model for revenue maximizing is different: We assume that bidders' values are drawn from publicly known probability distributions. However, we do not know the realizations, meaning the actual values. These are again private information.

1 Model

We again assume that there are n bidders; the set of all bidders is denoted by $\mathcal{N}$. Each of the bidders will report a bid b_i. We sell a single item among these bidders. Each bidder i has a private valuation $v_i \geq 0$ for being allocated the item. These values are drawn independently from publicly known distributions $\mathcal{D}_i$ of support $[0, v_{\max}]$. We assume that these distributions are continuous. Let the density function of $\mathcal{D}_i$ be denoted by f_i. Let the cumulative distribution function be denoted by F_i. That is,

$$F_i(t) = \int_{t'=0}^{t} f_i(t')\mathrm{d}t' = \mathbf{Pr}[v_i \leq t] \ .$$

We seek to design an allocation function $x\colon \mathbb{R}^n \to [0,1]^n$ that maps bids to probabilities of allocation with the constraint that $\sum_{i \in \mathcal{N}} x_i(b) \leq 1$. For today, we call this function x because f is used for the probability density. We pretend the function x is differentiable. The calculations remain correct although it is not.

Our main question today will be to find a *truthful* mechanism $\mathcal{M} = (x, p)$ that maximizes $\mathbf{E}_v \left[\sum_{i \in \mathcal{N}} p_i(v)\right]$ (among all truthful mechanisms).

That is, it is in each bidder's interest to tell the true value. Assuming that bidders tell us their true value, we want to maximize the revenue. This may sound a little strange: Why do we insist on truthfulness? We will come to this.

2 Example: One Item, One Bidder

If we have only a single bidder and one item, there is not a lot that we can do. By Myerson's lemma, the allocation has to be monotone in the bid. That is, there has to be a value p^* such that we sell the item if $b_1 \geq p^*$ and do not sell it otherwise. If $\mathcal{D}_1$ is the uniform distribution on $[0,1]$, then the expected revenue of any $p^* \in [0,1]$ is

$$\mathbf{E}_v \left[p_1(v)\right] = p^* \mathbf{Pr}\left[v_1 \geq p^*\right] = p^*(1 - p^*) \ ,$$

because we collect p^* if and only if the item is sold. This term is maximized for p^*.

3 Properties of the Revenue

Myerson's Lemma gives us a characterization what properties the functions x and p have to have. Namely, x has to be monotone and p follows the formula. These properties define the constraints of the optimization problem that we are solving, namely to find x and p so as to maximize $\mathbf{E}_v\left[\sum_{i\in\mathcal{N}} p_i(v)\right]$.

We first consider the payment of a single bidder keeping the other bids b_{-i} fixed. For a fixed value v_i, Myerson's Lemma tells us

$$p_i(v_i, b_{-i}) = \int_{t=0}^{v_i} t x_i'(t, b_{-i}) \mathrm{d}t \ .$$

Taking the expectation over v_i, we get

$$\mathbf{E}_{v_i}\left[p_i(v_i, b_{-i})\right] = \int_{v_i=0}^{v_{\max}} f_i(v_i) p_i(v_i, b_{-i}) \mathrm{d}v_i = \int_{v_i=0}^{v_{\max}} f_i(v_i) \int_{t=0}^{v_i} t x_i'(t, b_{-i}) \mathrm{d}t \mathrm{d}v_i \ .$$

Fubini's theorem tells us that we may switch the order of integration

$$\int_{v_i=0}^{v_{\max}} f_i(v_i) \int_{t=0}^{v_i} t x_i'(t, b_{-i}) \mathrm{d}t \mathrm{d}v_i = \int_{t=0}^{v_{\max}} \left(\int_{v_i=t}^{v_{\max}} f_i(v_i) \mathrm{d}v_i \right) t x_i'(t, b_{-i}) \mathrm{d}t = \int_{t=0}^{v_{\max}} (1-F_i(t)) t x_i'(t, b_{-i}) \mathrm{d}t$$

Now we do integration by parts: We differentiate $(1 - F_i(t))t$ and get $\frac{\mathrm{d}}{\mathrm{d}t}((1 - F_i(t))t) = -f_i(t)t + (1 - F_i(t))$. We integrate $x_i'(t, b_{-i})$, for which $\int x_i'(t, b_{-i}) \mathrm{d}t = x_i(t, b_{-i})$, so

$$\int (1 - F_i(t)) t x_i'(t, b_{-i}) \mathrm{d}t = (1 - F_i(t)) t x_i(t, b_{-i}) - \int \left(-f_i(t)t + (1 - F_i(t))\right) x_i(t, b_{-i}) \mathrm{d}t \ .$$

Overall this gives us

$$\begin{aligned}\int_{t=0}^{v_{\max}} (1 - F_i(t)) t x_i'(t, b_{-i}) \mathrm{d}t &= \underbrace{\left[(1 - F_i(t)) t x_i(t, b_{-i})\right]_{t=0}^{v_{\max}}}_{=0-0} - \int_{t=0}^{v_{\max}} \left(-f_i(t)t + (1 - F_i(t))\right) x_i(t, b_{-i}) \mathrm{d}t \\ &= \int_{t=0}^{v_{\max}} \left(f_i(t)t - (1 - F_i(t))\right) x_i(t, b_{-i}) \mathrm{d}t \ .\end{aligned}$$

We now define $\varphi_i(t) = t - \frac{1-F_i(t)}{f_i(t)}$ and rename t to v_i. This way

$$\mathbf{E}_{v_i}\left[p_i(v_i, b_{-i})\right] = \int_{v_i=0}^{v_{\max}} f_i(v_i) \varphi_i(v_i) x_i(v_i, b_{-i}) \mathrm{d}v_i = \mathbf{E}_{v_i}\left[\varphi_i(v_i) x_i(v_i, b_{-i})\right] \ .$$

Now, we include the other bidders by assuming $b_{-i} = v_{-i}$ (everybody bids truthfully) and taking the expectation over v_{-i}. Then we have

$$\mathbf{E}_v\left[p_i(v)\right] = \mathbf{E}_v\left[\varphi_i(v_i) x_i(v)\right] \ .$$

Taking the sum over all bidders and using linearity of expectation twice, we get

$$\mathbf{E}_v\left[\sum_{i\in\mathcal{N}} p_i(v)\right] = \sum_{i\in\mathcal{N}} \mathbf{E}_v\left[p_i(v)\right] = \sum_{i\in\mathcal{N}} \mathbf{E}_v\left[\varphi_i(v_i) x_i(v)\right] = \mathbf{E}_v\left[\sum_{i\in\mathcal{N}} \varphi_i(v_i) x_i(v)\right] \ .$$

We observe that this problem looks a lot like the problem of maximizing social welfare. In this case, we would have to find an allocation function x that maximizes $\sum_{i\in\mathcal{N}} v_i x_i(v)$. This we know is easy by selecting the bidder with the highest bid. The function $\sum_{i\in\mathcal{N}} \varphi_i(v_i) x_i(v)$ is called *virtual welfare* and each $\varphi_i(v_i)$ is called *virtual value.*

Lemma 21.1. *Let $\mathcal{M} = (x, p)$ be a truthful single-parameter mechanism, then the expected revenue equals the expected virtual welfare. That is,*

$$\mathbf{E}_v\left[\sum_{i\in\mathcal{N}} p_i(v)\right] = \mathbf{E}_v\left[\sum_{i\in\mathcal{N}} \varphi_i(v_i) x_i(v)\right], \ \textit{where } \varphi_i(t) = t - \frac{1 - F_i(t)}{f_i(t)} \ .$$

4 Regular Distributions

Lemma 21.1 tells us that maximizing the revenue is the same problem as maximizing the virtual welfare. There is one thing that we have to keep in mind: The allocation rule x has to be monotone in the bids. Therefore selecting the bidder with the highest (reported) virtual value is not always guaranteed to be monotone. If it is, then by charging payments according to the formula we get truthful mechanism.

The shape of the function φ_i depends on the distribution $\mathcal{D}_i$.

Definition 21.2. *A distribution $\mathcal{D}_i$ is* regular *if its associated virtual-value function φ_i is strictly increasing.*

You should be aware that the term *regular* is a little euphemistic. It is a reasonably strong assumption that often is not satisfied. Fortunately, however, there are enough positive examples.

Definition 21.3. *Define the* virtual-welfare maximizing mechanism *by allocation rule x^* that on input b maximizes $\sum_{i\in\mathcal{N}} \varphi_i(b_i)x_i^*$ and payments according to Myerson's lemma.*

Theorem 21.4. *If all bidders' distributions are regular, the virtual-welfare maximizing mechanism is truthful. Furthermore, it maximizes expected revenue among all truthful mechanisms.*

Proof. The allocation rule x^* is monotone if the distributions are regular. So it remains to show revenue optimality. Let $\mathcal{M} = (x, p)$ be an arbitrary truthful mechanism. Let p^* be the unique payment function according to Myerson's lemma that makes (x^*, p^*) truthful.

By Lemma 21.1, we have

$$\mathbf{E}_v\left[\sum_{i\in\mathcal{N}} p_i(v)\right] = \mathbf{E}_v\left[\sum_{i\in\mathcal{N}} \varphi_i(v_i)x_i(v)\right] \quad \text{and} \quad \mathbf{E}_v\left[\sum_{i\in\mathcal{N}} p_i^*(v)\right] = \mathbf{E}_v\left[\sum_{i\in\mathcal{N}} \varphi_i(v_i)x_i^*(v)\right] .$$

Furthermore $\varphi_i(v_i)x_i^*(v) \geq \varphi_i(v_i)x_i(v)$ for any v by the definition of x^*. Taking the expectation on both sides, this implies $\mathbf{E}_v\left[\sum_{i\in\mathcal{N}} p_i^*(v)\right] \geq \mathbf{E}_v\left[\sum_{i\in\mathcal{N}} p_i(v)\right]$. □

Now, what does it mean to maximize virtual welfare? If we sell a single item, the answer is surprisingly simple.

Lemma 21.5. *In the case of a single-item auction, virtual welfare is maximized by assigning the item to the bidder of the highest reported virtual value if this value is positive. Otherwise the item is left unallocated.*

Proof. Consider any valuation profile v and any allocation $x(v)$. If $\max_{i\in\mathcal{N}} \varphi_i(v_i) \geq 0$ and $\sum_{i\in\mathcal{N}} x_i(v) \leq 1$, then

$$\sum_{i\in\mathcal{N}} \varphi_i(v_i)x_i(v) \leq \max_{i\in\mathcal{N}} \varphi_i(v_i) \enspace ,$$

which is exactly the virtual welfare of the above allocation rule. If $\max_{i\in\mathcal{N}} \varphi_i(v_i) \leq 0$, then

$$\sum_{i\in\mathcal{N}} \varphi_i(v_i)x_i(v) \leq 0 \enspace ,$$

which is also the virtual welfare of the above allocation rule. So, in either case, the allocation has maximum virtual welfare. □

Example 21.6. *Consider the case that v_1 is drawn from the uniform distribution on $[0, 1]$, v_2 is drawn from the uniform distribution on $[0, 2]$. This way*

$$f_1(v_1) = 1 \qquad F_1(v_1) = v_1 \qquad \varphi_1(v_1) = v_1 - \frac{1 - v_1}{1} = 2v_1 - 1 \qquad \textit{for } v_1 \in [0, 1]$$

$$f_2(v_2) = \frac{1}{2} \qquad F_2(v_2) = \frac{1}{2}v_2 \qquad \varphi_2(v_2) = v_2 - \frac{1 - \frac{v_2}{2}}{\frac{1}{2}} = 2v_2 - 2 \qquad \textit{for } v_2 \in [0, 2]$$

If for example $v_1 = \frac{3}{4}$ *and* $v_2 = 1$*, then* $\varphi_1(v_1) = \frac{1}{2}$ *and* $\varphi_2(v_2) = 0$*. That is, bidder* 1 *wins the item. He has to pay the smallest value* t *for which he is a winner. In this case* $t = \frac{1}{2}$*.*

If $v_1 = \frac{1}{3}$ *and* $v_2 = \frac{2}{3}$ *then* $\varphi_1(v_1) = -\frac{1}{3}$ *and* $\varphi_2(v_2) = -\frac{2}{3}$*. Because both virtual values are negative, nobody gets the item.*