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 $\mathcal{D}_i$. 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,

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

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 task is to find a mechanism $\mathcal{M} = (x, p)$. We want it to be truthful and to maximize $\mathbf{E}_v\left[\sum_{i \in \mathcal{N}} p_i(v)\right]$. 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? The reason is that this assumption is to some extent without loss of generality. Intuitively, we could simulate whatever outcome a non-truthful mechanism has by a truthful one by acting strategically on a bidder's behalf. We will not formalize this "revelation principle" in this class.

2 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

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

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

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

Now we do integration by parts: We differentiate $(1 - F(t))t$ 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

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

$$= \int_{t=0}^{v_{\max}} \left(f_i(t)t - (1 - F_i(t))\right) x_i(t, b_{-i}) \mathrm{d}t$$

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

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

3 Regular Distributions

Lemma 11.1 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. However, this allocation rule.

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

Definition 11.2. *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 euphemistic. It is a reasonably strong assumption that often is not satisfied. Fortunately, however, there are enough positive examples.

Theorem 11.3. *If all bidders' distributions are regular, the revenue-maximizing auction assigns the item to the bidder if the highest reported virtual value if this value is positive, otherwise it leaves the item unallocated.*

Proof. Let $\mathcal{M} = (x, p)$ be an arbitrary truthful mechanism, and let x^* assign the item to the bidder if the highest reported virtual value if this value is positive, otherwise it leaves the item unallocated. Observe that this allocation rule maximizes virtual (declared) welfare, i.e., $\sum_{i \in \mathcal{N}} \varphi_i(v_i) x_i^*(v)$. To see this, we use that $\sum_{i \in \mathcal{N}} x_i^* \leq 1$. Furthermore, observe that if $\varphi_i(v_i) < 0$ for all i then $x_i^*(v) = 0$ for all i is maximizes the virtual welfare. If $\varphi_i(v_i) \geq 0$, then moving all mass to the highest entry is optimal.

Because the distributions are regular, the virtual-value functions φ_i are strictly increasing. This makes x^* monotone. Let p^* be the unique payment function according to Myerson's lemma that makes (x^*, p^*) truthful.

By Lemma 11.1, 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 $\mathbf{E}_v \left[\sum_{i \in \mathcal{N}} p_i^*(v) \right] \geq \mathbf{E}_v \left[\sum_{i \in \mathcal{N}} p_i(v) \right]$. □

Example 11.4. *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*

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

$$f_2(v_2) = \frac{1}{2} \qquad F_2(v_2) = \frac{1}{2} v_2 \qquad \varphi_2(v_2) = v_2 - \frac{1 - \frac{v_2}{2}}{\frac{1}{2}} = 2v_1 - 2 \qquad \textit{for } v_2 \in [0, 2]$$

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.

4 Identical Distributions

An interesting and enlightening special case is that all distributions are regular and identical. This way, the virtual-value functions φ_i become identical as well. Call it φ for now. The revenue-maximizing auction again gives the item the bidder i who maximizes $\varphi(b_i)$ if $\varphi(b_i) \geq 0$. Because of monotonicity and because the functions are identical, this is the bidder with maximal b_i. Call this bidder i^*.

The payment is again the smallest bid t that would make him a winner. Two possible cases can happen: If $\varphi(\max_{i \neq i^*} b_i) \geq 0$, then some other bidder would have won in the absence of i^*. So, he has to pay $t = \max_{i \neq i^*} b_i$. If $\varphi(\max_{i \neq i^*} b_i) < 0$, then nobody would have won the item. However, i^* still would have to bid so that $\varphi(t) \geq 0$. In other words, he has to pay $\varphi^{-1}(0)$.

Summarizing, the payment of bidder i^* is

$$\max \left\{ \varphi^{-1}(0), \max_{i \neq i^*} b_i \right\} .$$

That is, we have a second-price auction with a reserve price of $\varphi^{-1}(0)$. Just add a bidder bidding $\varphi^{-1}(0)$ to the second-price auction and if this bidder wins, nobody gets the item.

5 Beyond Truthfulness

Our insights today only concern truthful direct mechanisms. One would be tempted to think that non-truthful mechanisms might yield a higher revenue. For example, given the same bids, in a first-price auction the revenue is higher than in a second-price auction, which is truthful. However, this is not true if one takes into consideration the strategic behavior of the bidders.

On a high level, the argument works as follows. Given a non-truthful mechanism, the bidders will always choose a bid that depends on their current value. This means, bids follow bidding functions $v_i \mapsto b_i(v_i)$. As a result of strategic behavior, these bidding functions correspond to an equilibrium (the equilibrium concept is called *Bayes-Nash equilibrium*).

Instead of running this non-truthful mechanism, we can as well create a truthful mechanism that for each bidder first evaluates the bidding function b_i and then runs the non-truthful one. The outcome will be the same but bidders cannot improve by misreporting (otherwise we wouldn't have started from an equilibrium).

Proving these results formally requires introducing a new equilibrium concept (*Bayes-Nash equilibrium*), truthfulness concept (*Bayesian incentive compatibility*) but the key techniques are exactly the one we have seen today.