Revenue-Maximizing Auctions

Instructor: Xiaohui Bei

So far all mechanisms we studied focus on optimizing a specific objective function: the social welfare objective, which is defined as the sum of valuations of all agents. In practice, there are a number of other important objectives that people are interested in besides social welfare. Revenue, defined as the total amount of money that the auctioneer collects from the mechanism, is one obvious example. In this lecture we discuss auctions that are explicitly designed to maximize revenue.

1 Difficulties in Prior-Free Setting

As an example to illustrate the difficulties that one might encounter when the objective is maximizing revenue. Let us consider the simplest case in auctions, where there is only one item and one bidder with private valuation v for the item. According to the incentive compatibility characterization for the single-parameter environment, every incentive compatible mechanism has to be "posted price" mechanisms. That is, the auctioneer posts a fixed price of r. If the bidder has valuation $v \geq r$, the auction sells the item and generates a revenue of r. If $v < r$, the item is not sold and the revenue is 0.

Maximizing social welfare in this case is trivial: just set the price $r = 0$. Note that this is an optimal strategy *independent* of the value v. However, when comes to the revenue objective, one can immediately realize that there does not exist a mechanism that can even approximately optimize the revenue for all v. For example, a mechanism with posted price r will do terribly if on inputs smaller than r or much bigger than r. Therefore, in order to have any meaningful solutions to this problem, one has to be provided with some additional information about the private valuation v.

2 Bayesian Analysis

To deal with this issue, in this lecture we introduce the most classical and well-studied setting: the *Bayesian* analysis. We again focus on the domain of single-parameter settings. In a Bayesian setting, we assume the following additional information:

- Each bidder i's valuation v_i is independently drawn from a distribution F_i with density function f_i and support within the interval $[0, v_{\max}]$. Recall that $F_i(x)$ denotes the probability that a random variable drawn from F has value at most x.
- The distributions $F_1, \dots, F_n$ are known in advance to the auctioneer. The realizations $v_1, \dots, v_n$ of bidders are private information as before.
- The goal for the auctioneer is to design an incentive compatible and individual rational auction that maximizes his expected revenue, where the expectation is with respect to the given distribution $F_1 \times F_2 \times \cdots \times F_n$. Such auction is also called an *optimal auction.*

Remark. It is worth noting that for the analysis in this lecture, we assume that distributions $F_1, \dots, F_n$ are unknown information to each bidder. Hence we focus on incentive compatibility in expectation where the expectation is over the randomness of the mechanism, but not the distributions $F_1, \dots, F_n$. Although the main results in today's lecture also apply more generally to "Bayes-Nash incentive compatible" auctions, in which case the bidders must know the distributions $F_1, \ldots, F_n$.

2.1 One Buyer One Item Revisited

With these additional information, now let us reconsider the problem of selling one item to one buyer. This time, we further assume that the bidder's valuation v is drawn from distribution F. Again one considers mechanisms with posted price r. Now the expected revenue is simply $r \cdot [1 - F(r)]$. With different F, one can compute the optimal price r that maximize this term. The resulting auction will be the optimal auction for this setting. For example, if F is the uniform distribution on $[0, 1]$, the optimal price is $1/2$, and in turn the optimal expected revenue will be $1/4$.

One can further generalize this example to two bidders. For example, if both v_1 and v_2 are uniformly distributed on $[0, 1]$. The Vickrey auction has the revenue equals to the expected value of the smaller bid, which is $1/3$. Then we could also supplement the Vickrey auction with a *reserved price*: a Vickrey auction with reserved price r gives the item to the highest bidder, unless all bids are less than r. And the payment is the higher value between r and the second highest bid. One can do the computation and derive that with two bidders, the optimal reserved price is $1/2$, and the resulting auction has an expected revenue of $5/12$, which is better than the original Vickrey auction. But we are not sure if this is the optimal auction in this setting. There might be other formats of incentive compatible auctions that have better revenues. The analysis becomes much more complicated due to the increase of space of incentive compatible mechanisms. In the rest of this lecture, we will talk about Myerson's theorem, which provides a complete solution to this problem in the single-parameter setting.

3 Virtual Welfare

Recall that in general single-parameter setting with n bidders and one item, a mechanism consists of an allocation function $x_i(\mathbf{b})$ and a payment function $p_i(\mathbf{b})$ for each bidder. The expected revenue of this mechanism is

$$\mathbb{E}_{\mathbf{b}}\left[\sum_i p_i(\mathbf{b})\right] = \sum_i \left[\int_{\mathbf{b}_{-i}} f_{-i}(\mathbf{b}_{-i}) \left(\int_{b_i} f_i(b_i) p_i(\mathbf{b}) \,\mathrm{d}b_i\right) \mathrm{d}\mathbf{b}_{-i}\right]$$

Right now we do not know how to analyze this objective. But one thing we do know how to analyze is the social welfare:

$$\mathbb{E}_{\mathbf{b}}\left[\sum_i v_i(\mathbf{b})\right] = \sum_i \left[\int_{\mathbf{b}_{-i}} f_{-i}(\mathbf{b}_{-i}) \left(\int_{b_i} f_i(b_i) v_i x_i(\mathbf{b}) \,\mathrm{d}b_i\right) \mathrm{d}\mathbf{b}_{-i}\right]$$

These two formulas look very much alike. Hence we wonder: can we reduce expected revenue maximization to expected social welfare maximization in some "virtual" space?

That is, we hope to find some virtual valuation function $\phi_i(b_i)$ for each bidder i. Such that the expected revenue of a mechanism $\mathbb{E}[\sum_i p_i(\mathbf{b})]$ is always equal to the expected "virtual social welfare" $\mathbb{E}[\sum_i \phi_i(b_i) \cdot x_i(\mathbf{b})]$.

In order for this to hold, it suffices for the virtual valuations to satisfy

$$\int_0^{v_{\max}} f_i(b_i) p_i(b_i) \,\mathrm{d}b_i = \int_0^{v_{\max}} f_i(b_i) \phi_i(b_i) x_i(b_i) \,\mathrm{d}b_i$$

for every i and $\mathbf{b}_{-i}$.

By Myerson's Lemma, we can replace the payment function $p_i(b_i)$ by $b_i x_i(b_i) - \int_0^{b_i} x_i(t)\,\mathrm{d}t$. Then the expected payment of bidder i becomes

$$\begin{aligned}
\mathbb{E}[p_i(b_i)] &= \int_0^{v_{\max}} f_i(b_i) p_i(b_i)\,\mathrm{d}b_i \\
&= \int_0^{v_{\max}} f_i(b_i) \left(b_i x_i(b_i) - \int_0^{b_i} x_i(t)\,\mathrm{d}t \right) \mathrm{d}b_i \\
&= \int_0^{v_{\max}} b_i x_i(b_i) f_i(b_i)\,\mathrm{d}b_i - \int_{b_i} \int_0^{b_i} x_i(t) f_i(b_i)\,\mathrm{d}t\,\mathrm{d}b_i
\end{aligned}$$

Next we switch the order of integration in the second term, and then rename t to b_i to get

$$\begin{aligned}
\mathbb{E}[p_i(b_i)] &= \int_0^{v_{\max}} b_i x_i(b_i) f_i(b_i)\,\mathrm{d}b_i - \int_0^{v_{\max}} x_i(t) \int_t^{v_{\max}} f_i(b_i)\,\mathrm{d}b_i\,\mathrm{d}t \\
&= \int_0^{v_{\max}} b_i x_i(b_i) f_i(b_i)\,\mathrm{d}b_i - \int_0^{v_{\max}} x_i(t)\,(1 - F_i(t))\,\mathrm{d}t \\
&= \int_0^{v_{\max}} b_i x_i(b_i) f_i(b_i)\,\mathrm{d}b_i - \int_0^{v_{\max}} x_i(b_i)\,(1 - F_i(b_i))\,\mathrm{d}b_i \\
&= \int_0^{v_{\max}} \left[b_i - \frac{1 - F_i(b_i)}{f_i(b_i)} \right] x_i(b_i) f_i(b_i)\,\mathrm{d}b_i
\end{aligned}$$

Hence in order to let $\mathbb{E}[p_i(b_i)] = \int_0^{v_{\max}} f_i(b_i)\phi_i(b_i)x_i(b_i)\,\mathrm{d}b_i$, we can just let

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

Such definition of $\phi_i(v_i)$ is called the *virtual valuation* of bidder i with valuation v_i. Note that this virtual valuation depends on bidder i's own valuation and his distribution, and not on those of others. And we have just showed that in every auction,

Theorem 11.1. *The expected revenue of any incentive compatible mechanism, is equal to its expected virtual welfare, i.e.,* $\mathbb{E}[\sum_i p_i(\mathbf{b})] = \mathbb{E}[\sum_i \phi_i(b_i) \cdot x_i(\mathbf{b})]$.

Hence we can now reduce the problem of revenue maximization to the problem of welfare maximization.

4 Bayesian Optimal Auctions

The above theorem allows us to describe the following Myerson's optimal mechanism:

(1) Given bids $\mathbf{b}$ and distributions $F_1, \ldots, F_n$, compute "virtual bids" $b'_i = \phi_i(b_i)$ for each bidder i.

(2) Run VCG mechanism on the virtual bids $\mathbf{b}'$ to get the allocation rule $\mathbf{x}'$ and payments $\mathbf{p}'$.

(3) Output $(x, p) \leftarrow (x', \phi^{-1}(p_i))$.

It seems that Myerson's optimal mechanism can help us maximize the virtual welfare, hence revenue, of the mechanism. However, there is still one issue left: is this mechanism incentive compatible? By the incentive compatibility characterization theorem, this depends on whether or not we have a monotone allocation rule. Hence our question is: is this virtual welfare-maximizing rule monotone? The answer to this question depends on the distribution F. In fact, if we virtual valuation function ϕ is monotone increasing, it can be easily seen that the virtual welfare-maximizing allocation rule must also be monotone.

Definition 11.2. *A distribution F is* regular *if the corresponding virtual valuation function $\phi(v) = v - \frac{1-F(v)}{f(v)}$ is monotone increasing.*

Lemma 11.3. *Myerson's optimal mechanism is incentive compatible if and only if for every bidder i, distribution F_i is regular.*

Myerson's work can also be extend to accommodate the case where ϕ is not monotone. But we do not cover it here.

4.1 Examples Revisited

Now let us return to the single-item auction with n i.i.d. bidders. We further assume that for each bidder the distribution is regular. Hence the optimal auction allocates the item to the bidder with the highest virtual valuation. Except that there is a catch: in general we assume the valuation of a bidder is always nonnegative, but even so the virtual valuation could still be negative. When the highest virtual valuation is negative, the virtual welfare-maximizing allocation should not allocate this item to any bidder.

When bidders are all i.i.d, they have the same virtual valuation. This allocation rule is equivalent to the Vickrey auction with a reserved price of $\phi^{-1}(0)$.

Corollary 11.4. *When bidders are i.i.d with regular distributions, the optimal auction is Vickrey auction with reserved price $\phi^{-1}(0)$.*

For example, if every bidder i has an uniform distribution on $[0, 1]$, his virtual valuation function becomes $\phi(v_i) = 2v_i - 1$. In this case, the optimal auction is just the Vickrey auction with reserved price 1/2.

Recommended Literature

- Chapter 13.1, 13.2 in the AGT book.
- J. D. Hartline. Mechanism design and approximation. Book draft. September, 2014.
- Tim Roughgarden's lecture notes `http://theory.stanford.edu/~tim/f13/l/l5.pdf` and lecture video `https://youtu.be/jQsAoMcxlIo`
- R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58-73, 1981.