

18.1 Introduction

Today we're going to finally talk about an obvious question: maximizing revenue. After all, over the last few weeks we've been taking on the role of the auctioneer. So maximizing revenue is a very natural goal.

Unfortunately, with our current setup there's basically nothing we can do. To see this, consider a basically trivial case: single-item single-bidder. That is, there's only a single bidder, who has a private valuation v for the one item that's being sold. Maximizing social welfare with an incentive compatible auction is trivial: give the item to the bidder at price 0. In fact, this is almost the only thing we can do: clearly in this situation the only incentive compatible mechanism is a "posted-price" mechanism: we decide on a price r before anything happens, and then the single bidder comes along and decides whether or not to buy at that price. So if we want to have an IC mechanism that maximizes revenue, all we can do is try to "guess" the true valuation v and set $r \approx v$. If we guess correctly then we're approximately maximizing welfare, but if $r \ll v$ or $r > v$ then we're out of luck.

So in order to talk about revenue-maximizing auctions, we need to change the setup by making extra assumptions. We're basically going to take a "average-case" or "Bayesian" perspective on things. We'll talk a little bit later about whether this is reasonable, but it's at least a place to start. And, perhaps most importantly, it leads to some really interesting results and connections back to what we've been talking about (social surplus).

18.2 Setup and Examples

- Single parameter environment with players $[n]$ and feasible set $X \subseteq \mathbb{R}^n$.
- For each $i \in [n]$ there is a distribution F_i .
 - We'll assume that each F_i has support in $[0, v_{\max}]$ for some $v_{\max}$.
 - Let $F_i(z) = \Pr_{x \sim F_i}[x \leq z]$ denote the cumulative distribution function (CDF) of F_i .
 - Let f_i be the probability density function, so $\int_0^z f_i(z) dz = F_i(z)$.
 - These distributions are known to the mechanism.
 - Bidder i has some private $v_i \sim F_i$.
- We're looking for the "optimal" mechanism: maximum expected revenue among all incentive compatible mechanisms, where the expectation is over valuations drawn from the distributions.

18.2.1 Examples

Single-item Single-Bidder. Consider the setup we started with. Now we have a known distribution F for the valuation of the single player. This means that if we set the price to r , our expected revenue is $r(1 - F(r))$ (the revenue from making a sale times the probability of making the sale). So we just need to find the r which maximizes this. For example, if F is uniform on $[0, 1]$, then we are trying to maximize $r(1 - r)$, so the best choice is to set $r = 1/2$, in which case we get expected revenue of $1/4$.

Single-item Two-Bidder. Suppose there are two bidders instead of one, and suppose that F_1 and F_2 are both uniform on $[0, 1]$. One option for an auction is the second-price (Vickrey) auction, since we know that this is incentive compatible. In this auction, the expected revenue we make is the expectation of the second highest valuation, which is $1/3$.

Is there a better IC auction? From Myerson, if we use the allocation rule of “give to the highest bidder” then the second-price auction is our only choice. But can we actually do better by having a different allocation rule?

Consider the “reserve price r ” auction: our allocation rule is that we give the item to the highest bidder *if* that bidder bids at least r . If the top bid is less than r , then no one gets the item. This might seem like a crazy allocation rule, since sometimes we’re forcing ourselves to get 0 revenue by refusing to sell. But it turns out that it actually does better. By Myerson, the price we charge (to the winner) to make this mechanism incentive compatible is $\max(r, \text{second highest bid})$. If we set $r = 1/2$, then our expected revenue is

$$\begin{aligned} & \Pr[\text{both values above } 1/2] \cdot \mathbf{E}[\text{second highest value} \mid \text{both values above } 1/2] \\ & + \Pr[\text{highest value above } 1/2, \text{second value below } 1/2] \cdot \frac{1}{2} \\ & = \frac{1}{4} \cdot \frac{2}{3} + \frac{1}{2} \cdot \frac{1}{2} = \frac{2}{12} + \frac{3}{12} = \frac{5}{12} > 1/3 \end{aligned}$$

So we can design an incentive-compatible auction that’s better than the second-price auction! Can we do even better? Maybe by setting a different reserve price, or maybe even a totally different format?

18.3 Virtual Welfare and Revenue-Optimal Auctions

Let’s start trying to answer our main question: in our setting (single-parameter environment with valuations drawn from known distributions), what is the revenue-maximizing incentive compatible mechanism? Since we’re in a single-parameter environment and only care about incentive compatible mechanisms, Myerson applies. So if we figure out a monotone allocation rule $\mathbf{x} : [0, v_{\max}]^n \rightarrow X$, the prices $\mathbf{p}$ that we charge are determined by Myerson. We’ll let $p_i(b) = \mathbf{p}(b)_i$.

So this means that the revenue-maximizing IC auction is the monotone allocation rule which maximizes

$$\mathbf{E}_{v \sim F_1 \times F_2 \times \dots \times F_n} \left[\sum_{i=1}^n p_i(v) \right].$$

So in some sense we've solved the problem, but this is super unsatisfying. What is this allocation rule? Does it have any structure? Can we compute it?

Definition 18.3.1. *The virtual valuation of player i with valuation v_i is*

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

So the virtual valuation depends on the true valuation and on the distribution from which it was drawn. One thing that's important to note: the virtual valuation, unlike the true valuation, could be negative. For example, if F_i is uniform between $[0, 1]$, then $F_i(v_i) = v_i$ and $f_i(v_i) = 1$. So then

$$\phi_i(v_i) = v_i - \frac{1 - v_i}{1} = 2v_i - 1$$

The main theorem of today is the following.

Theorem 18.3.2. *Let $\mathbf{x}$ be a monotone allocation rule with payment function $\mathbf{p}$ given by Myerson's Lemma. Then*

$$\mathbf{E}_{v \sim F_1 \times F_2 \times \dots \times F_n} \left[\sum_{i=1}^n p_i(v) \right] = \mathbf{E}_{v \sim F_1 \times F_2 \times \dots \times F_n} \left[\sum_{i=1}^n \phi_i(v_i) x_i(v) \right].$$

So the expected revenue is equal to the expected *virtual* welfare (surplus)! We'll try to prove this by the end of the day, but it's basically just a bunch of calculus. Right now, let's see how we can use this to find optimal auctions. We've connected revenue maximization with surplus maximization, which is something that we know quite a bit about.

With Theorem 18.3.2 in hand, the obvious mechanism is the virtual welfare-maximizing allocation rule:

$$\mathbf{x}(v) = \arg \max_{x \in X} \left(\sum_{i=1}^n \phi_i(v_i) x_i(v) \right)$$

This clearly maximizes the expected virtual welfare, and so by Theorem 18.3.2 the expected revenue, but only if it is monotone. It turns out that this depends on the distributions: there are distributions where the fact that we use virtual valuations make this allocation rule non-monotone and thus non-implementable.

Definition 18.3.3. *A distribution F is regular if its virtual valuation function $v - \frac{1 - F(v)}{f(v)}$ is non-decreasing in v .*

Many natural distributions are regular, including the uniform distribution, exponential distributions, and log normal distributions. But others aren't, including some multimodal distributions and some heavy-tailed distributions.

Theorem 18.3.4. *If F_i is regular for all $i \in [n]$, then the virtual welfare-maximizing allocation rule (with consistent tie-breaking) is monotone.*

As an immediate corollary, we get that if all players have valuations from regular distributions then the revenue-maximizing incentive compatible mechanism is precisely the virtual welfare-maximizing allocation rule together with prices from Myerson.

Proof Sketch of Theorem 18.3.4. This is essentially the same as the proof for knapsack auctions from last lecture (which we saw applied to the welfare-maximizing allocation in any single-parameter environment), so I'm not really going to do it. The key point is that since all the distributions are regular, higher bids result in higher virtual values. Thus bidding more results in getting at least as much stuff allocated to you, since you appear to be even more valuable to the virtual welfare-maximizing allocation. $\square$

So now we know optimal auctions when values are from regular distributions! This is still a bit unsatisfying, though – is there any simple description or intuition about these auctions?

Example. Suppose we're selling a single item, there are n bidders, and each bidder has a valuation drawn i.i.d. from the *same* regular distribution F (for simplicity, let's assume that F is strictly regular, i.e., the virtual valuation function is increasing rather than just nondecreasing). Then the virtual welfare maximizing allocation is the allocation maximizing $\sum_{i=1}^n \phi(v_i)x_i$ over all $x \in \{0, 1\}^n$ such that $\sum_{i=1}^n x_i \leq 1$, where ϕ is the virtual valuation function for F . What allocation is this? Since F is strictly regular, the bidder with maximum $\phi(v_i)$ is the bidder with maximum v_i , i.e., $\arg \max_{i \in [n]} \phi(v_i) = \arg \max_{i \in [n]} v_i$. So it seems like we want allocate the item to the bidder with maximum valuation, just like in the Vickrey (second-price) auction!

There is an important difference, though: the virtual valuation might be negative! So if the maximum virtual valuation is negative, then the best allocation is to actually not give the item away. So the following allocation rule is actually the right one. Let i^* be the bidder with largest virtual valuation (and thus largest actual valuation). If $\phi(v_{i^*}) > 0$, allocate the item to i^* , and otherwise do not give the item to anyone. The payment for this auction are from Myerson: we charge the winner $\max(\phi^{-1}(0), \text{second highest bid})$.

In other words, this is the Vickrey auction with reserve price $\phi^{-1}(0)$! So, for example, if F is uniform between 0 and 1, then $\phi(v) = 2v - 1$ and thus we run a Vickrey auction with reserve price $1/2$.

18.4 Proof of Theorem 18.3.2

Fix some agent $i \in [n]$ and all other valuations v_{-i} (at this point these are arbitrary, but later we will require them to be drawn from the appropriate distributions). Let's assume that the allocation function x_i for player i is differentiable (to handle the more general case around discontinuities we would just need some fancier calculus). Then we know from Myerson that the price function for player i is

$$p_i(v) = \int_0^{v_i} z \cdot x'_i(v_{-i}, z) dz$$

where we're using x'_i to denote the derivative of x_i with respect to z . Now we can write a bunch of equalities.

$$\begin{aligned}
\mathbf{E}_{v_i \sim F_i} [p_i(v)] &= \int_0^{v_{\max}} p_i(v_{-i}, v_i) f_i(v_i) dv_i && \text{(def of expectation)} \\
&= \int_0^{v_{\max}} \left(\int_0^{v_i} z \cdot x'_i(v_{-i}, z) dz \right) f_i(v_i) dv_i && \text{(Myerson)} \\
&= \int_0^{v_{\max}} \left(\int_z^{v_{\max}} f_i(v_i) dv_i \right) \cdot z \cdot x'_i(v_{-i}, z) dz && \text{(reverse order of integration)} \\
&= \int_0^{v_{\max}} (1 - F_i(z)) \cdot z \cdot x'_i(v_{-i}, z) dz && \text{(def of } F_i \text{ and } f_i)
\end{aligned}$$

Now we can do integration by parts to proceed. If, like me, you've totally forgotten this, a quick reminder: integration by parts says that $\int f g' dx = f g - \int g f' dx$, where f and g are functions of x . So let's look at the formula we have and set things appropriately: we'll use integration by parts where $f(z) = (1 - F_i(z)) \cdot z$ and where $g'(z) = x'_i(v_{-i}, z)$. Then $g(z) = x_i(v_{-i}, z)$ and $f'(z) = 1 - F_i(z) + z(-f_i(z)) = 1 - F_i(z) - z f_i(z)$ (where we applied the product rule and the fact that the derivative of F_i is f_i). So now when we apply integration by parts, we get

$$\begin{aligned}
\mathbf{E}_{v_i \sim F_i} [p_i(v)] &= [(1 - F_i(z)) z \cdot x_i(v_{-i}, z)]_0^{v_{\max}} - \int_0^{v_{\max}} x_i(v_{-i}, z) (1 - F_i(z) - z f_i(z)) dz \\
&= 0 - 0 - \int_0^{v_{\max}} x_i(v_{-i}, z) (1 - F_i(z) - z f_i(z)) dz \\
&= - \int_0^{v_{\max}} x_i(v_{-i}, z) (1 - F_i(z) - z f_i(z)) dz.
\end{aligned}$$

Continuing, we get

$$\begin{aligned}
\mathbf{E}_{v_i \sim F_i} [p_i(v)] &= \int_0^{v_{\max}} \left(z - \frac{1 - F_i(z)}{f_i(z)} \right) x_i(v_{-i}, z) f_i(z) dz && \text{(factor out } -f_i(z)) \\
&= \int_0^{v_{\max}} \phi_i(z) x_i(v_{-i}, z) f_i(z) dz && \text{(def of } \phi_i) \\
&= \mathbf{E}_{z \sim F_i} [\phi_i(z) x_i(v_{-i}, z)] && \text{(def of } f_i \text{ and expectation)} \\
&= \mathbf{E}_{v_i \sim F_i} [\phi_i(v_i) x_i(v)]
\end{aligned}$$

Note that the above equality is true for every v_{-i} . So now we can just use linearity of expectations.

$$\begin{aligned}
\mathbf{E}_{v \sim F_1 \times F_2 \times \dots \times F_n} \left[\sum_{i=1}^n p_i(v) \right] &= \sum_{i=1}^n \mathbf{E}_{v \sim F_1 \times F_2 \times \dots \times F_n} [p_i(v)] \\
&= \sum_{i=1}^n \mathbf{E}_{v \sim F_1 \times F_2 \times \dots \times F_n} [\phi_i(v_i) x_i(v)] \\
&= \mathbf{E}_{v \sim F_1 \times F_2 \times \dots \times F_n} \left[\sum_{i=1}^n \phi_i(v_i) x_i(v) \right]
\end{aligned}$$