

19.1 Introduction

Last class we showed how to design truly optimal auctions, with respect to revenue maximization. But these had a few weaknesses: they could get quite complex, and they assumed that we had precise information about the distributions that each bidder's valuation was from. Today we're going to talk about some ways around these issues.

19.1.1 Review

Last time we saw that maximizing revenue is the same as maximizing “virtual” welfare. Let's think about the optimal auction from last class in the case where we're selling a single item. Recall that each player has a publicly known regular distribution F_i , and a private valuation $v_i \sim F_i$. Thanks to our theorem that maximizing expected revenue is the same thing as maximizing expected virtual welfare, we know that the optimal auction is to give the item to $\arg \max_{i \in [n]} \phi_i(v_i)$ if $\max_{i \in [n]} \phi_i(v_i) > 0$, where $\phi(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ is the virtual valuation.

If $F_i = F$ for all $i \in [n]$, i.e., all players have the same distribution, then the bidder with maximum virtual valuation is also the bidder with highest real valuation. So in this setting the auction is pretty simple: it's a Vickrey auction with reserve $\phi^{-1}(0)$.

19.2 Prior-Free Auctions

One issue with the optimal auctions from last class was the assumption that we precisely know the valuation distributions. Let's think about this even in the simplest setting, where we're selling a single item and all players have the same valuation distribution F . We know from last class that the optimal auction in this setting is the Vickrey auction with reserve price $\phi^{-1}(0)$, i.e., if the highest bidder has bid larger than $\phi^{-1}(0)$ then we give it to them at price $\max(\phi^{-1}(0), \text{second highest bid})$. This is pretty simple, but it requires that we precisely know F in order to set the reserve price correctly. What can we do if we don't know F at all?

The following theorem is due to Bulow and Klemperer, and gives a very interesting bound for this situation (single-item, multiple bidders with the same unknown F). For any regular distribution F , let OPT_F be the optimal auction (Vickrey with correct reserve) for F .

Theorem 19.2.1. *Let F be a regular distribution and let $n \in \mathbb{N}$. Then*

$$\mathbf{E}_{v_1, \dots, v_{n+1} \sim F} [\text{Revenue}(\text{Vickrey on } v_1, v_2, \dots, v_{n+1})] \geq \mathbf{E}_{v_1, \dots, v_n} [\text{Revenue}(OPT_F \text{ on } v_1, \dots, v_n)]$$

In other words, the revenue we get from our old second-price (Vickrey) auction with $n+1$ players is at least the revenue that we would get from the optimal auction with n players. The usual interpretation of this is that adding one more bidder is more important than getting the auction *exactly* right – if we need to invest resources to either learn/estimate F , we might as well just use

a second-price auction and use those resources to convince one more player to play.

Proof of Theorem 19.2.1. To prove this, we’re going to consider a “fake” auction which is neither Vickrey nor optimal, but which we will be able to relate both Vickrey and optimal to. This auction \mathcal{A} has $n+1$ bidders (with valuations drawn from F) and works as follows: it runs OPT_F on bidders in $[n]$, but if the item is not sold to one of these bidders (i.e., no one meets the reserve) then it gives the item to bidder $n+1$ for free.

It is easy to see that \mathcal{A} is incentive compatible, since if you’re one of the bidders in $[n]$ the auction is the exact same as OPT_F , and if you’re bidder $n+1$ then nothing you do affects the outcome and you’re never charged anything. Let’s first note two obvious but important properties of \mathcal{A} :

1. $\mathbf{E}[\text{Revenue}(\mathcal{A} \text{ on } n+1 \text{ bidders})] = \mathbf{E}[\text{Revenue}(OPT_F \text{ on } n \text{ bidders})]$. This is obvious because \mathcal{A} runs OPT_F on n bidders, and if it doesn’t sell the item there then it gets 0 extra revenue from giving the item to the $n+1$ bidder.
2. \mathcal{A} always sells the item.

Now we claim that the Vickrey auction maximizes revenue among all incentive compatible auctions that always sell the item. If we can prove this, then we’ve proved the theorem: it implies that Vickrey does at least as well as \mathcal{A} (by the second property), and we know that \mathcal{A} does exactly as well as OPT_F (by the first property). So we just need to prove this claim.

We know that expected revenue equals the expected virtual welfare. So if we are forced to always sell the item, the best we can do is give the item to the bidder with the highest virtual valuation, since this will maximize $\sum_{i=1}^{n+1} \phi(v_i)x_i$ over all allocations $x \in \{0,1\}^{n+1}$ such that $\sum_{i=1}^{n+1} x_i = 1$. Since F is regular, this means that we should give the item to the highest bidder, and then the price (by Myerson) is the second highest bid. So if we’re forced to sell the item, Vickrey is the best we can do. \square

19.3 Highest Bidder (Almost)

Let’s move to a slightly more general setting: still just selling one item, but now each bidder i has a possibly different distribution F_i .

19.3.1 Review + Setup

Let’s think about the optimal auction from last class in the case where we’re selling a single item. Recall that each player has a publicly known regular distribution F_i , and a private valuation $v_i \sim F_i$. Thanks to our theorem that maximizing expected revenue is the same thing as maximizing expected virtual welfare, we know that the optimal auction is to give the item to $\arg \max_{i \in [n]} \phi_i(v_i)$ if $\max_{i \in [n]} \phi_i(v_i) > 0$, where $\phi(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ is the virtual valuation. The price we charge is from Myerson, and is the “critical bid” needed to secure the item: if player i gets the item, we charge them $\max(\phi_i^{-1}(0), \phi_i^{-1}(\max_{j \neq i} \phi_j(v_j)))$.

If the F_i distributions are different, though, then this auction is a little weird. The top bidder might not win (since the bidder with the largest valuation is not necessarily the bidder with the

highest virtual valuation), and the prices are hard to explain: they have to do with both ϕ_i and ϕ_j and so take into account the distributions of all the players. So, in particular, to justify these prices to some player we need to assume that the player knows and understands not just their own distribution, but all of the other distributions as well.

So we have a somewhat natural question: are there “simpler” auctions which, while not optimal, are at least almost-optimal? Note that we don’t mean “simple” in a formal running-time sense, but rather in a more informal sense.

19.3.2 The Mechanism

One natural way of trying to make the auction simpler is to force ourselves to either give the item to the top bidder or to no one. This at least makes the outcome more “intuitive” and “simpler” to the bidders: the highest bidder is the one who gets the item, just like we would naturally expect in a single-item auction. So let’s try to design such an auction in a way which also has relatively simple prices. We’re not quite going to be able to make it this simple, but we’ll get pretty close.

To simplify notation a bit, for $z \in \mathbb{R}$ we will let $z^+ = \max(0, z)$.

Our prices are (as usual) forced by Myerson, so the question is what allocation rule to use. We clearly need to sometimes give the item to no one, or else we’re just doing the Vickrey auction precisely. We’re going to use an allocation rule that is slightly more complicated than just Vickrey + reserve (like what we had with one distribution) but which is still relatively simple: Vickrey with bidder-specific reserves.

To formalize this, let $t \in \mathbb{R}_{\geq 0}$ such that $\Pr[\max_{i \in [n]} \phi_i(v_i)^+ \geq t] = 1/2$ (if no such t exists, for example if the distributions are not continuous, then there are standard workarounds that we’re not going to get into – see Exercise 6.2 from Roughgarden’s book). We’ll set a reserve price of $r_i = \phi_i^{-1}(t)$ for each player i (in other words, the virtual valuation has to be at least t). Then given a bid vector, we give the item to the highest bidder who has met their reserve (if no one meets their reserve then we don’t sell the item). It’s not hard to see that if all the F_i distributions are regular then this allocation rule is monotone (good exercise to do at home), and then by Myerson the price that we charge the winner i is $\max(r_i, \max_{j \neq i: v_j \geq r_j} v_j)$.

This is a pretty intuitive auction to explain to the bidders: each bidder has a reserve price, and we give it to the highest bidder who has met their reserve at either their reserve or the second-highest bid (among those bidders who have met their reserve).

Theorem 19.3.1. *If all valuation distributions are regular, then this auction has expected revenue at least 1/2 the expected revenue of the optimal auction.*

To prove this theorem, we’re going to have to take a small diversion to something called a *prophet inequality*. This is actually slightly overkill for what we need, but it’s beautiful and will also be important later when we talk about online auctions. So, like when we took a diversion to learning theory to analyze no-regret dynamics, we’re going to put auctions on hold for a minute.

19.3.3 Prophet Inequality

This is a problem in online algorithms / stopping theory. There are n distributions G_1, G_2, \dots, G_n , all of which we know. At time i , we are offered a “prize” $\pi_i \sim G_i$. We can either accept this prize and end the process, or reject it and move to the next time. Our goal is to maximize our expected reward (value of the prize we accept). In particular, our goal is to get expected reward that is close to $\mathbf{E}[\max_{i \in [n]} \pi_i]$, which is the expected reward of a “prophet” who knows all of the prizes at the beginning of time. Can we design an algorithm that gets almost as much?

Consider the following algorithm: we set a threshold of t and accept prize i if and only if $\pi_i \geq t$. This gives a family of *threshold strategies*. We will set $t \in \mathbb{R}$ to be the value such that $\Pr[\max_{i \in [n]} \pi_i \geq t] = 1/2$.

Theorem 19.3.2. *The expected value of this threshold strategy is at least $\frac{1}{2} \mathbf{E}[\max_{i \in [n]} \pi_i]$.*

Proof. Let’s first provide some bounds on the threshold strategy, and then we’ll provide a comparable bound on the prophet.

By the definition of t , with probability $1/2$ we don’t get any prize at all. On the other hand, with probability $1/2$ we get reward at least t . In fact, we might get more. If π_i is the *only* prize above t , then we will definitely choose prize i and so will get π_i reward, i.e., an *additional* $\pi_i - t$. If there are multiple prizes above t , then this is trickier – how much additional reward we get depends on which of these prizes come first. So let’s just be pessimistic and say that if there are multiple prizes above t , we only get t . Then

$$\begin{aligned} \mathbf{E}[\text{reward}] &\geq \frac{1}{2}t + \sum_{i=1}^n \mathbf{E}[\pi_i - t | (\pi_i \geq t) \wedge (\pi_j < t \ \forall j \neq i)] \cdot \Pr[(\pi_i \geq t) \wedge (\pi_j < t \ \forall j \neq i)] \\ &= \frac{1}{2}t + \sum_{i=1}^n \mathbf{E}[\pi_i - t | \pi_i \geq t] \cdot \Pr[\pi_i \geq t] \cdot \Pr[\pi_j < t \ \forall j \neq i] \\ &\geq \frac{1}{2}t + \sum_{i=1}^n \mathbf{E}[\pi_i - t | \pi_i \geq t] \cdot \Pr[\pi_i \geq t] \cdot \frac{1}{2} \\ &= \frac{1}{2}t + \frac{1}{2} \sum_{i=1}^n \mathbf{E}[(\pi_i - t)^+] \end{aligned}$$

Now let's think about the prophet.

$$\begin{aligned}
\mathbf{E} \left[\max_{i \in [n]} \pi_i \right] &= \mathbf{E} \left[t + \max_{i \in [n]} (\pi_i - t) \right] \\
&= t + \mathbf{E} \left[\max_{i \in [n]} (\pi_i - t) \right] \\
&\leq t + \mathbf{E} \left[\max_{i \in [n]} (\pi_i - t)^+ \right] \\
&\leq t + \mathbf{E} \left[\sum_{i=1}^n (\pi_i - t)^+ \right] \\
&= t + \sum_{i=1}^n \mathbf{E}[(\pi_i - t)^+]
\end{aligned}$$

Thus the expected reward of the algorithm is at least half the expected reward of the prophet, as claimed. \square

19.3.4 Auction Analysis

Now let's use our analysis of the prophet inequality to analyze the Vickrey with bidder-specific reserves auction that we discussed. This is pretty straightforward: just think of $\phi_i(v_i)^+$ as the prize π_i . Then the bidder-specific reserves just translates into saying that we take the highest bid of all the bidders whose prize is at least t , and this is the same t as in the prophet inequality. So the auction is the exact same as the prophet inequality, except that the auction takes the highest bid (with prize above the threshold) while the prophet inequality takes the first prize (above the threshold). But if you look at our analysis, if there was more than one prize above the threshold we only gave ourselves the threshold value anyway! So this doesn't actually make any difference in the analysis.

Slightly more formally, note that our mechanism chooses some prize $\phi_i(v_i)^+$ as long as at least one prize is above t . The prophet inequality says that if we take the *first* prize above t , then the expected value is at least $\frac{1}{2} \mathbf{E}[\max_{i \in [n]} \phi_i(v_i)^+]$. But the analysis of the theorem worked as long as we took an *arbitrary* prize above t : it didn't matter that we took the first. Thus the theorem also applies to taking the *max bid* (which might not be the max prize) as long as the associated prize is above the threshold.

Thus if x is the allocation vector of our mechanism, we have that

$$\mathbf{E} \left[\sum_{i=1}^n \phi_i(v_i)^+ x_i \right] \geq \frac{1}{2} \mathbf{E} \left[\max_{i \in [n]} \phi_i(v_i)^+ \right].$$

Now since $t \geq 0$, our mechanism only sells the item if there is a bidder i with $\phi_i(v_i) \geq 0$. Thus

$$\mathbf{E} \left[\sum_{i=1}^n \phi_i(v_i)^+ x_i \right] = \mathbf{E} \left[\sum_{i=1}^n \phi_i(v_i) x_i \right],$$

so our allocation has expected virtual welfare at least $1/2$ the expected virtual welfare of the optimal auction. Since revenue equals virtual welfare, the same guarantee is true of expected revenue.