6.853 Algorithmic Game Theory and Data Science March 19, 2019

Lecture 12

Lecturer: Vasilis *Scribe: Devendra*

1 Introduction

In the last lecture, we learned about Bayesian single-dimensional settings for which Myerson's theorem states that there exists a revenue-optimal auction that is direct and DSIC. Furthermore, the expected optimal revenue is equal to the sum of expected virtual welfare of individual bidders, which depends on the bidder distributions.

However, the auctioneer may not know the distributions. In this lecture, we will look at auctions that work without the knowledge of distributions, and yet compete as well as the auctions as if they know everything about the distributions.

1.1 Complexity of Myerson's auction

Under single-item setting, with regular independent bidders, when the distributions of bidders are the same, i.e. $F_1 = F_2 \cdots = F_n$, then the optimal auction takes a simple form: allocate to the bidder with the highest value and charge the maximum of the second highest bid and some reserve price. This is the well-known second-price auction with a reserve price. This fact follows from the observation that Myerson's optimal auction allocates to the player with the highest virtual value. However, when all virtual value functions are monotone (by regularity) and identical (by i.i.d. assumption), then this is equivalent to allocating to the player with the highest value (subject to his virtual value being positive). The threshold above which the virtual value is positive, designates the value of the optimal reserve price for the distribution. Moreover, the second price payment rule makes the latter allocation rule truthful.

However, when $F_1 \neq F_2 \neq \cdots \neq F_n$, even if the distributions are regular and hence all $\phi()$'s are monotone, the optimal Myerson auction is no longer so simple as is evident by the following example.

Example 1. *Consider the case of two bidders and one item. The value of bidder* 1 *is drawn from* $v_1 \sim U[0,1]$ *and the value of bidder* 2 *is drawn from* $v_2 \sim U[0,100]$*. In this case,* $\phi_1(v_1) = 2v_1 - 1$ *and* $\phi(v_2) = 2v_2 - 100$*. The allocation rule of the optimal auction in this case is as follows:*

- *When* $v_1 > \frac{1}{2}, v_2 < 50$*, allocate to 1 and charge* $\frac{1}{2}$*.*
- *When* $v_1 < \frac{1}{2}, v_2 > 50$*, allocate to 2 and charge* 50*.*
- *When* $0 < 2v_1 - 1 < 2v_2 - 100$*, allocate to 2 and charge:* $(99 + 2v_1)/2$*, a tiny bit above 50.*
- *When* $0 < 2v_2 - 100 < 2v_1 - 1$*, allocate to 1 and charge* $(2v_2 - 99)/2$*, a tiny bit above* $1/2$*.*

The above characterization stems from the fact that the allocation rule of the optimal aucton must be maximizing virtual welfare. Notice that, in this setting, the auction allocates the item to bidder 1, even when its value is, say 0.75, and other bidder 2's value is, say 40, which is substantially greater than 0.75.

In non i.i.d. single-dimensional settings, Myerson's auction is hard to explain to people without the knowledge of virtual valuations. And, the optimal revenue auction would inevitably weird. This motivates the topic of today's lecture. Are there simpler, more practical, and more robust auctions than the theoretically optimal auction? Optimality requires complexity. Therefore, we will look for approximately optimal, but simple auctions.

2 Simple Auction: Second Price with Player-Specific Reserves

The first simple auction rule that we will investigate is a small generalization of the second price auction with a reserve price. Namely consider the following auction:

- For each bidder i associate a reserve price r_i
- Among all bidders with $v_i \geq r_i$, allocate to the highest value player
- If the item is allocated to some player i, then charge the maximum of the next highest bid and his reserve price r_i

The only change that we made to the second price auction with a reserve price (which was the optimal auction in the i.i.d. setting) was that we allowed a different reserve price for each bidder. This different reserve price allows us to price discriminate and use the prior information that we have about the value of each bidder to extract more revenue (i.e. we can set a higher reserve price to bidders which we know that their value is stochastically hgiher).

The main question: is this degree of price discrimination enough to guarantee a good fraction of the optimal revenue? or do we need the fully complex allocation rule that Myerson's optimal auction designates? We will see that the answer is that player-specific reserve prices suffice! In particular, we will show that they always achieve at least half of the optimal revenue. In fact we will see that this can be achieved by an even simpler version of the auction described above, where we allocate to any player that passes his reserve price (even uniformly at random). To do this we will do a small detour to a very relevant problem in optimal stopping theory.

2.1 Optimal Stopping Rules

Consider the following game:

- There are n stages.
- In stage i, you are offered a nonnegative prize Π_i drawn from distribution G_i.
- You are given the distributions $G_1, \cdots, G_n$ before teh game begins, and told that the prizes are drawn independently from these distributions.
- But each Π_i is revealed at the beginning of stage i.
- After seeing Π_i, you can either accept the prize and end the game, or discard the prize and proceed to the next stage.

Question: Is there a strategy for playing the game, whose expected reward competes with that of a prophet who knows all realized Π_is and picks the largest?

The difficulty in answering this question stems from the trade-off between the risk of accepting a reasonable prize and missing out on a better one later vs. the risk of having to settle for a lousy prize in one of the final stages.

Theorem 1 (Prophet Inequality)**.** *Suppose, we are presented with n prizes $\Pi_1, \cdots, \Pi_n$ sequentially, each rewards Π_i drawn from distribution G_i. There exists a stopping time τ such that*

$$\mathbb{E}[\Pi_\tau] \geq \frac{1}{2}\mathbb{E}[\max_i \Pi_i]$$

In fact a threshold strategy: pick first $\Pi_i \geq \theta$ works well.

Proof. Let's be very generous with OPT. If everything is below θ, then it collects at most θ. If there is some rewards above θ, then it collects $(\max_i \Pi_i - \theta)_+$ maximum excess payoff above θ.

It at most, always collects $\sum_i (\Pi_i - \theta)_+$ the excess payoff above θ of every prize, i.e.

$$\Pi_\star = \max_i \Pi_i = \theta + (\Pi_\star - \theta)_+ \leq \theta + \sum_i (\Pi_i - \theta)_+$$

where $a_+ = max(a, 0)$. Therefore,

$$OPT \leq \theta + \sum_i \mathbb{E} \underbrace{[(\Pi_i - \theta)_+]}_{A_i}$$

where $a_+ = max(a, 0)$. □

What about APX? It collects θ only when $\Pi_\star \geq \theta$ (i.e. there exists one reward above θ). Then it also collects some excess payoff: $\mathbb{E}[(\Pi_\tau - \theta)_+]$. This excess payoff is equal to $\mathbb{E}[(\Pi_i - \theta)_+ | \Pi_i \geq \theta]$ when { all rewards for $j \neq i$ are below θ, except i }. These are disjoint events. So,

$$\begin{aligned} \text{Excess} = & \sum_i \mathbb{E}[(\Pi_i - \theta)_+ | \Pi_i \geq \theta] P(\Pi_i \geq \theta) P(\forall j \neq i : \Pi_j < \theta) \\ = & \sum_i \mathbb{E}[(\Pi_i - \theta)_+] P(\forall j \neq i : \Pi_j < \theta) \end{aligned}$$

The probability that $\{\forall j \neq i : \Pi_j < \theta\}$ is at least $P(\forall j : \Pi_j < \theta) = P(\Pi_\star < \theta)$.

$$\text{Excess} = P(\Pi_\star < \theta) \sum_i A_i$$

$$\begin{aligned} APX = \mathbb{E}[\Pi_\tau \mathbf{1}\{\Pi_\tau \geq \theta\}] &= \mathbb{E}[\theta \mathbf{1}\{\Pi_\tau \geq \theta\}] + \mathbb{E}[(\Pi_\tau - \theta)\mathbf{1}\{\Pi_\tau \geq \theta\}] \\ &= \mathbb{E}[\theta \mathbf{1}\{\Pi_\tau \geq \theta\}] + \mathbb{E}[(\Pi_\tau - \theta)_+] \\ &= \theta P(\Pi_\star \geq \theta) + \sum_i \mathbb{E}[(\Pi_i - \theta)_+ | \tau = i] P(\tau = i) \\ &\geq \theta P(\Pi_\star \geq \theta) + \sum_i \mathbb{E}[(\Pi_i - \theta)_+ | \Pi_i \geq \theta, \Pi_j < \theta \ \forall \ j \neq i] P(\Pi_i \geq \theta, \Pi_j < \theta \ \forall \ j \neq i) \\ &= \theta P(\Pi_\star \geq \theta) + \sum_i \mathbb{E}[(\Pi_i - \theta)_+ | \Pi_i \geq \theta] \cdot P(\Pi_i \geq \theta) \cdot P(\Pi_j < \theta \ \forall \ j \neq i) \\ & \qquad \text{(independence of } \Pi_i\text{)} \\ &\geq \theta P(\Pi_\star \geq \theta) + \sum_i A_i \cdot P(\forall \ j \neq i : \Pi_j < \theta) \\ &\geq \theta P(\Pi_\star \geq \theta) + \sum_i A_i \cdot P(\forall \ j : \Pi_j < \theta) \\ &= \underbrace{\theta P(\Pi_\star \geq \theta) + P(\Pi_\star < \theta)}_{\text{Pick } \Pi_\star \text{ s.t. } P(\Pi_\star \geq \theta) = \frac{1}{2}} \sum_i A_i \\ &= \frac{1}{2}(\theta + \sum_i A_i) \end{aligned}$$

One drawback of the latter analysis is that it assumes that there exists a point θ such that $\Pr[\Pi_* \geq \theta] = 1/2$. The latter might not always be the case for discontinuous prize distributions. However, we could overcome this difficulty with an alternative choice of θ (albeit a bit less intuitive). Another way to define the threshold θ is to choose θ so that the two terms that constitute the bound on the optimal revenue are equal, i.e. θ, such that:

$$\theta = \sum_i \mathbb{E}[(\Pi_i - \theta)_+] \tag{1}$$

Observe that such a point θ always exists, since the two quantities (LHS and RHS) are continuous functions in θ. Moreover, the LHS is increasing in θ for $\theta \geq 0$ and the second quantity is decreasing in θ for $\theta \geq 0$. Moreover, the LHS is 0 at 0 and the RHS is non-negative at $\theta = 0$. Thus there must exist a solution $\theta \geq 0$ to the above equation. For such a theta, observe that: $OPT \leq 2\theta$ and that $APX = \theta \Pr[\Pi_* \geq \theta] + \Pr[\Pi_* < \theta]\theta = \theta \geq OPT/2$.

2.2 Revisiting the Revenue of the Simple Auction

Observe that we can analyze the revenue of the second price auction with player specific reserves as an optimal stopping rule problem. We can view each bidder as a prize with value equal to the positive part of their virtual value $\Pi_i = \phi_i^+(v_i) = \max\{\phi_i(v_i), 0\}$. Then the revenue of the Myerson optimal auction is equal to the prize of the optimal stopping rule, i.e. $\mathbb{E}[\max_i \Pi_i] = \mathbb{E}[\max_i \phi_i^+(v_i)]$.

However, we know that there is a simple threshold strategy that collects at least half of the optimal prize, i.e. at least half of the expected optimal virtual welfare. Hence, at least half of the optimal revenue. This threshold strategy essentially says: if I choose a threshold θ such that:

$$\Pr[\max_i \phi_i^+(v_i) \geq \theta] = 1/2 \tag{2}$$

Then we know that allocating to any player that satisfies $\phi_i^+(v_i) \geq \theta$, guarantees at least half of the optimal revenue. Since ϕ_i are monotone, the latter is equivalent to allocating to any player with value $v_i \geq r_i$ for some r_i. The second price auction with player-specific reserves is a particular allocaton function that satisfies this condition. We could have even allocated uniformly at random to any player that passes his reserve price.

3 Prior-independent Auctions

Another critique to the optimal auction is that it assumes the bidder distributions. This assumption can be considered valid if there exists sufficient data. However, that may not be the case in general, especially in the case when the market is "thin". Also, the auctioneer may not be confident about the bidders' distributions. These considerations motivate us to a different class of auctions known as prior-independent auctions. Specifically, can there be auctions that do not use any knowledge of distributions, and yet perform nearly as well as if they knew everything about the distributions.

Theorem 2 (Bulow-Klemperer '96)**.** *Consider any regular distribution F and n:*

$$\mathbb{E}_{v_1,\cdots,v_{n+1}\sim F}[REV(\textit{Vickrey})] \geq \mathbb{E}_{v_1,\cdots,v_{n+1}\sim F}[REV(\textit{Myerson})]$$

Proof. Consider another auction M with $n+1$ bidders:

1. First, we run the Myerson auction on the first n bidders.
2. If the item is unallocated to the first n bidders, then, we allocate it to the last bidder for free.

This is a DSIC mechanism. Moreover, it has the same revenue as Myerson's auction with n bidders. Note that, the allocation rule always allocates the item to some bidder.

Vickrey auction also always allocates the item to the bidder with highest value (also with the highest virtual value).

Vicrey Auction has the highest virtual welfare among all DSIC mechanisms that always allocate the item. □

Remark 1. *1. Vickrey auction is prior-independent.*

2. The theorem implies that more competition is better than finding the right auction format.

Corollary 1. *If we had one sample r of the distribution:*

$$\mathbb{E}_{v_1,\cdots,v_n\sim F}[REV(\textit{Vickrey with reserve } r)] \geq \left(1 - \frac{1}{n+1}\right) \mathbb{E}_{v_1,\cdots,v_n\sim F}[REV(\textit{Myerson})]$$

Proof. Observe that we can view a Vickrey auction among n bidders with a random reserve r drawn from the same distribution F as the bidder values as follows: we introduce an $n+1$ bidder to the auction whose value is equal to the reserve price. We run a Vickrey auction among these $n+1$ bidders. If the item is allocated to this new fake bidder we simply don't allocate the item and don't collect any revenue. Observe that by the symmetry of the Vickrey auction, the expected revenue contribution of every player is equal to $Rev/(n+1)$, where Rev is the expected revenue of the auction. Thus by discarding the revenue of the fake bidder we lose at mot $1/(n+1)$ of the revenue of the Vickrey auction among $(n+1)$ bidders, i.e.:

$$\mathbb{E}_{v_1,\cdots,v_n\sim F}[\text{REV(Vickrey with reserve } r)] \geq \left(1-\frac{1}{n+1}\right)\mathbb{E}_{v_1,\cdots,v_n,v_{n+1}\sim F}[\text{REV(Vickrey)}] \tag{3}$$

Invoking the Bulow-Klemperer theorem, we also get that the RHS of the above equation is at least $\left(1-\frac{1}{n+1}\right)\mathbb{E}_{v_1,\cdots,v_n\sim F}[\text{REV(Myerson)}]$, which concludes the proof. $\square$

We can also similarly show another corollary, that does not require, neither a random sample r from the distribution, nor an extra bidder. Albeit the following corollary gives a meaningful bound only when $n \geq 2$:

Corollary 2.

$$\mathbb{E}_{v_1,\cdots,v_n\sim F}[\mathit{REV(Vickrey)}] \geq \left(1-\frac{1}{n}\right)\mathbb{E}_{v_1,\cdots,v_n\sim F}[\mathit{REV(Myerson)}]$$

Proof. Invoking the Bulow-Klemperer theorem for $n-1$, we have:

$$\mathbb{E}_{v_1,\cdots,v_n\sim F}[\text{REV(Vickrey)}] \geq \mathbb{E}_{v_1,\cdots,v_{n-1}\sim F}[\text{REV(Myerson)}],$$

Thus it suffices to show the optimal revenue with $n-1$ bidders is at least $1-1/n$ of the optimal revenue with n bidders.

We will do this as follows: consider the following way of generating a random sample of n i.i.d. valuations $v_1,\ldots,v_n$ from distribution F

- Sample n numbers $u_1,\ldots,u_n$ from F
- Sample a random permutation π of these n numbers
- Set $v_i = u_{\pi(i)}$

Let $\vec{u}=(u_1,\ldots,u_n)$ and $\vec{v}=(v_1,\ldots,v_n)$. Now observe that:

$$OPT(n) := \mathbb{E}_{v_1,\cdots,v_n\sim F}[\text{REV(Myerson)}] = \mathbb{E}_{\vec{v}\sim F^n}\left[\max_{i\in\{1,\ldots,n\}}\phi(v_i)\right] = \mathbb{E}_{\vec{u}\sim F^n,\pi}\left[\max_{i\in\{1,\ldots,n\}}\phi(u_{\pi(i)})\right]$$

Similarly:

$$OPT(n-1) := \mathbb{E}_{v_1,\cdots,v_{n-1}\sim F}[\text{REV(Myerson)}] = \mathbb{E}_{\vec{u}\sim F^n,\pi}\left[\max_{i\in\{1,\ldots,n-1\}}\phi(u_{\pi(i)})\right]$$

Conditional on the random vector $\vec{u}$, $OPT(n)$ is different from $OPT(n-1)$, only if the highest $\phi(u_i)$ happens to fall in the last slot of the permutation π. Since the permutation was uniformly random, the latter happens with probability $1/n$. In that event $OPT(n-1)$ is smaller than $OPT(n)$ by at most $OPT(n)$. Thus: $OPT(n) - OPT(n-1) \leq OPT(n)/n$, which yields the corollary. $\square$