6.853 Algorithmic Game Theory and Data Science	March 14, 2019
Lecture 11	
Lecturer: Constantinos Daskalakis	*Scribe: Govind & Devendra*

1 Introduction

In the last lecture, we learned about single-dimensional settings which concern with bidders (players) each of whom have a scalar private value. We learned the definition of implementable and monotone allocation rules. Then we presented Myerson's Lemma which states that:

- An allocation rule is implementable if and only if it is monotone.
- For an implementable/monotone rule there is an *essentially unique* payment function such that the sealed-bid mechanism is DSIC.
- In particular, there is a *unique* payment function such that the mechanism is DSIC, IR with non-positive transfers.

Finally, we saw some applications of Myerson's Lemma.

In this lecture, we will try and find allocation and payment rules such that mechanism results in maximum revenue for the auctioneer.

2 Single Item Revenue Maximization

Previously we looked at maximizing social welfare in single item auctions. Today we will look at the trickier question of revenue maximization.

First off, let us consider the following thought experiment. Suppose there is one bidder whose value for the item is v, and that is the highest value. Then if we were to give a take-it-or-leave-it offer, the highest achievable revenue is v. Unfortunately, the value v is *private*, so how can we possibly price the item optimally?

This simple issue is actually quite fundamental. If we fix *any* price, there is some v for which it fails to extract revenue v. This indicates that the bidder's private value is simply too strong of a benchmark for the auction to complete against.

We will solve this by assuming that the distribution F over the bidder's value is known, and perform an average-case analysis over the distribution.

Example 1. *Let us consider a simple case with one bidder and one item. If we post price r, we get expected revenue $r \cdot (1 - F(r))$. So if F is the uniform distribution on $[0, 1]$, the optimal choice of r is $1/2$, which achieves expected revenue $1/4$. The optimal posted price is also called the* monopoly price. *Can we do better? Namely, in the same scenario with one bidder and a uniform value distribution over $[0, 1]$, is there any auction that does better than just posting the price 1/2 up front?*

Example 2. *Now suppose there are instead two bidders whose values are uniformly distributed over $[0, 1]$. The revenue of Vickrey's auction is the expectation of the minimum of the two uniform random variables, which gives us $1/3$. Is there anything we can do to do better?*

Perhaps surprisingly, the answer is yes! By including a reserve *price of 1/2, we get expected revenue $5/12 > 1/3$[1]. Is 5/12 the optimal expected revenue of* any *auction?*

[1] Recall that a reserve price is a minimum price under which the seller refuses to sell the good.

3 Revenue-Optimal Auctions

We will be interested in Bayesian, single-dimensional environments. In these, we have n bidders, where each bidder i has a private (scalar) value v_i, which is her value for being served, and v_i is sampled from the distribution F_i, where F_i is known to everyone (the auctioneer and other bidders). Finally, there is a feasible set X of n-dimensional vectors $(x_1, \ldots, x_n)$ where x_i denotes whether bidder i is served. Examples of these include k-unit auctions and sponsored search.

When we are in this setting, it turns out that there are simple, direct, DSIC, and IR *revenue-optimal* auctions!

Theorem 1 (Myerson '81)**.** *Consider a single-dimensional setting, where the distribution F_i of every bidder's private value is known to all other bidders as well as the auctioneer. Then there exists a* revenue-optimal *auction that is simple, direct, DSIC, and individually rational.*

Myerson's auction is *optimal*, in the sense that the expected revenue when all (which is in their best interest, since the auction is DSIC) is as large as the expected revenue of any other (potentially indirect) auction, when bidders use Bayesian Nash equilibrium strategies. In particular, the revenue is as large as that of any BIC direct mechanism.

Theorem 2. *Fix a Bayesian single-dimensional environment, where bidder distributions are $F_1, \ldots, F_n$ and $F = F_1 \times \ldots \times F_n$. Let (x, p) be a BIC mechanism satisfying interim IR and NPT. The expected revenue of this mechanism under truth-telling is*

$$\mathbb{E}_{v\sim F}[\sum_i p_i(v)] = \mathbb{E}_{v\sim F}[x_i(v)\varphi_i(v_i)] \tag{1}$$

where $\varphi_i(v_i) := v_i - (1 - F_i(v_i))/f_i(v_i)$ is bidder i's "virtual value function" (and f_i denotes the density function for F_i).

In particular, note that while the LHS of Equation 1 is the expected revenue, the RHS looks like an expected welfare, if the values are $\varphi_i(v_i)$. Hence, Equation 1 says that the expected revenue is equal to the expected virtual welfare, which we now prove.

Proof.

$$\begin{aligned}
\mathbb{E}_{v\sim F}[\sum_i p_i(v)] &= \sum_i \mathbb{E}_{v\sim F}[v_i \cdot x_i(v_i, v_{-i}) - \int_0^{v_i} x_i(t, v_{-i})dt] \\
&= \sum_i \mathbb{E}_{v_i \sum F_i}[v_i \mathbb{E}_{v_{-i}}[x_i(v_i, v_{-i})] - \int_0^{v_i} \mathbb{E}_{v_{-i}}[x_i(t, v_{-i})]dt] \\
&= \sum_i \mathbb{E}_{v_i\sim F_i}[v_i \hat{x}_i(v_i) - \int_0^{v_i} \hat{x}_i(t)dt]
\end{aligned}$$

where $\hat{x}$ is the interim allocation to bidder i. Now we bring in our knowledge of the distribution of v_i:

$$
\begin{aligned}
&= \sum_i \mathbb{E}_{v_i \sim F_i}[v_i \hat{x}_i(v_i)] - \sum_i \int_{v_i=0}^{+\infty} \int_{t=0}^{v_i} \hat{x}_i(t) f_i(v_i) dt dv_i \\
&= \sum_i \mathbb{E}_{v_i \sim F_i}[v_i \hat{x}_i(v_i)] - \sum_i \int_{t=0}^{+\infty} \int_{v_i=t}^{+\infty} \hat{x}_i(t) f_i(v_i) dt dv_i \\
&= \sum_i \mathbb{E}_{v_i \sim F_i}[v_i \hat{x}_i(v_i)] - \sum_i \int_{t=0}^{+\infty} \hat{x}_i(t)(1 - F_i(t)) dt \\
&= \sum_i \int_{v_i=0}^{+\infty} v_i \cdot \hat{x}_i(v_i) f(v_i) dv_i - \sum_i \int_{0}^{+\infty} \hat{x}_i(v_i)(1 - F_i(v_i)) dv_i \\
&= \sum_i \int_0^{+\infty} \hat{x}_i(v_i) \cdot \left(v_i - \frac{1 - F_i(v_i)}{f(v_i)} \right) f(v_i) dv_i \\
&= \sum_i \mathbb{E}_{v_i}[\hat{x}_i(v_i) \cdot \varphi_i(v_i)] \\
&= \mathbb{E}_{v \sim F}[\sum_i x_i(v) \cdot \varphi_i(v_i)]
\end{aligned}
$$

□

4 Interim allocation and rule

In a Bayesian, single-dimensional setting, given a direct auction (x, p), we define its interim allocation rule and interim price rule for bidder i as follows:

$$
\begin{aligned}
\hat{x}_i(v_i) &= \mathbb{E}_{v_{-i}}[x_i(v_i; v_{-i}] \\
\hat{p}_i(v_i) &= \mathbb{E}_{v_{-i}}[p_i(v_i; v_{-i}]
\end{aligned}
$$

The interim allocation rule expresses the allocation that bidder i expects to receive when reporting v_i in expectation over the other bidders' values, assuming that the report truthfully. Similarly, for the interim price rule.

For notational simplicity, we often omit the hats, representing both the interim and the ex-post allocation rule for bidder i by $x_{i_{(\cdot)}}$ with the meaning of the function being determined by the numer of arguments.

Similarly, we denote both the interim and ex-post price rule by $p_i(\cdot)$.

The Myerson's Lemma for Bayesian setting is as follows:

- An allocation rule x is *BIC* implementable if and only if it is *interim* monotone (i.e. for all $i, x_i(v_i)$ is point-wise increasing in v_i).
- If x is BIC implementable, there is an *essentially unique interim* payment function such that the direct sealed-bid mechanism (x, p) is BIC, given by the formula:

$$
\forall\ i, v_i : \mathbb{E}_{v_{-i}}[p_i(v)] = \mathbb{E}_{v_{-i}} \left[v_i x_i(v) - \int_0^{v_i} x_i(t; v_{-i}) dt + p_i(0, v_{-i}) \right] \quad (2)
$$

- In particular, there is a *unique interim* payment rule such that the mechanism is BIC and additionally interim IR with non-positive transfers.

Remark 1. *The proof of this lemma is essentially identical to its DSIC version from Lecture 10.*

Example 3. *Recall the one U[0,1] bidder, one item example. We saw that posting aprice of* $1/2$ *achieves expected revenue of* $1/4$*. We now show that* $1/4$ *is the optimal expected revenue.*

Proof. *The virtual transform for* $U[1,0]$ *distribution is* $\phi(v) = v - \frac{1-F(v)}{f(v)} = 2v - 1$*. By Myerson's theorem, any BIC, IR, NPT mechanism* (x, p) *has revenue of* $\mathbb{E}_{v\sim U[0,1]}[x(v).(2v-1)] \leq \frac{1}{4}$*. Hence, posting a price of* $1/2$ *is optimal as it achieves the best possible expected virtual welfare, and hence expected revenue.* □

Example 4. *Similarly, we prove for two U[0,1] bidders, one item case, that the optimal expected revenue is* $5/12$ *which is achieved using Vickrey auction with reserve.*

Proof. *The virtual transform for* $v \sim U[0,1]$ *is* $\phi(v) = 2v - 1$*.*

By Myerson's theorem, optimizing expected revenue is equivalent to optimizing expected virtual welfare subject to interim monotonicity (needed for BIC).

On a bid profile (v_1, v_2)*, the pair of virtual values are* $(\phi(v_1), \phi(v_2)) = (2v_1-1, 2v_2-1)$*. The allocation rule for under non-positive transfers would rule out allocation if the maximum of the virtual values is less than 0.* $max(v_1, v_2) \leq 2$*. So, we should only consider the case* $max(v_1, v_2) \geq 2$*. In this case, the item should be allocated to the bidder with highest virtual value.*

The above allocation rule is monotone, so by Myerson's Lemma ther is a price rule that makes it DSIC.

Now, an allocation rule which is DSIC and pointwise optimal virtual welfare is also revenue-optimal.

Finally, we note that the above auction is identical to Vickrey auction with reserve $1/2$*, which yields a revenue of* $5/12$ *and is also DSIC.* □

5 Myerson's Optimal Revenue Auction

We now proceed to characterize an optimal revenue auction which results from Myersons' theorem.

Note that by revelation principle, instead of optimizing over all direct/indirect interim IR auctions, we can limit ourselves to optimizing over direct interim IR, BIC auctions. Moreover, one can also assume interim NPT because, if interim NPT is not satisfied, then expected revenue can be increased by adding a term to make it interim NPT. Hence, by Myerson's theorem, it suffices to find an *interim monotone* allocation rule that optimizes expected *virtual welfare*.

Well, what if we only limit attention to determining an allocation rule which optimizes expected virtual welfare? The answer is to identify an allocation that maximizes virtual welfare point-wise, i.e. for every bid profile v, sove: $\max \sum_i x_i \phi(v_i)$ s.t. $x \in X$. Let us call this "Virtual Welfare-Maximizing Rule".

If the Virtual Welfare-Maximizing Rule is monotone, then we have a revenue-optimal single-dimensional auction which is also DSIC!

However, the monotonicity depends on the type of distribution. For example, if a bidder has value with a distribution $2/3U[0,1] + 1/3U[0,1]$, then in the auction for allocating one item, the virtual welfare-maximizing rule is not monotone.

5.1 Regular distributions

Definition 1 (Regular distributions)**.** *A single-dimensional distribution* F *is* regular *if its virtual transform* $v - \frac{1-F(v)}{f(v)}$ *is non-decreasing.*

Definition 2 (Monotone Hazard Rate)**.** *A single-dimensional distribution* F *has Monotone Hazard Rate, if* $\frac{1-F(v)}{f(v)}$ *is non-increasing.*

- MHR distributions include uniform, exponential, Gaussian distributions, etc.
- MHR and Power-law are regular distributions.

- Multi-modal or distributions with heavy tails are irregular distributions.

Theorem 3. *Fix a Bayesian single-dimensional environment, where all bidders draw values from regular distributions.*

Then the auction whose allocation rule is the virtual welfare maximizing allocation rule (and whose price is uniquely determined by the allocation rule so that the resulting auction is interim IR, NPT) is DSIC and revenue-optimal (among all interim IR, potentially inndirect auctions).