COMP/MATH 533: Algorithmic Game Theory	October 1, 2014

Lecture 9

Lecturer: Yang Cai — *Scribe: Rory Bokser and Joey Litalien*

NOTE: The content of these notes has not been formally reviewed by the lecturer. It is recommended that they are read critically.

1 Myerson's Auction Recap

1.1 Myerson's Theorem

Recall the statement of Myerson's Theorem:

Theorem (Myerson, '81). *For any single-dimensional environment, let $\mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2 \times \cdots \times \mathcal{F}_n$ be the joint value distribution, and $(\mathbf{x}, \mathbf{p})$ be a DSIC mechanism. The expected revenue of this is mechanism is given by*

$$\mathbf{E}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} p_i(\mathbf{v})\right] = \mathbf{E}_{\mathbf{v}\sim\mathcal{F}}\left[\sum_{i=1}^{n} x_i(\mathbf{v})\phi_i(v_i)\right],$$

where $\phi_i(v_i) := v_i - (1 - F_i(v_i))/f_i(v_i)$ is called bidder i's virtual value (f_i is the density function for $\mathcal{F}_i$).

Myerson showed that using the above result, if we want to maximize the revenue, a simple auction can do the trick. Such auction, called *Myerson's Auction*, are done as follows:

- Bidders report their values;
- The reported values are *transformed* into *virtual values*;
- The virtual-welfare maximizing allocation is chosen;
- Charge the payments according to Myerson's Lemma;
- Transformation depends on the distributions $\mathcal{F}_i$'s; it is a deterministic function (the virtual value function).

1.2 Some nice properties of Myerson's Auction

- **DSIC**, but optimal among all Bayesian Incentive Compatible (BIC) mechanisms.
- **Deterministic**, but optimal among all possibly randomized mechanisms. You might think that a randomized mechanism would give higher revenue, yet Myerson's auction is maximal. If we are in a single-dimensional setting, the best mechanism is always deterministic. This does *not* hold for multi-dimensional settings as we shall see below.[1]
- **Central open problem** in mathematical economics: How can we extend Myerson's result to multi-dimensional settings? Important progress has been made in the past few years. We will not discuss these in class.

[1] Recall from Lecture 3 that a *single dimensional environment* is defined as having n bidders, with each bidder i having a private valuation v_i representing its value "per unit of stuff". There is also a feasible set X, each element of which is an n-dimensional vector $\mathbf{x} = (x_1, \ldots, x_n)$, where x_i denotes the "amount of stuff" given to bidder i.

2 Challenges in Multi-Dimensional Settings

Besides being computationally more challenging, the optimal auctions in multi-dimensional environments are also structurally more involved. We will consider three examples with two items and a single bidder, and we will see that many of the counterintuitive phenomena already exist in such simple cases.

2.1 First example

In this example, we will assume that the bidder is *additive*, that is, $v(\{1,2\}) = v_1 + v_2$. To further simplify the setting, we will also suppose that v_1 and v_2 are both drawn i.i.d from a distribution $\mathcal{F} = U\{1,2\}$, i.e. $\Pr[v_i = 1] = \frac{1}{2} = \Pr[v_i = 2]$.

What is the optimal auction here? It is somewhat unclear. The natural attempt would be to sell both items using Myerson's Auction separately. Intuitively, we would think that this gives a good revenue. It is clear that in this case, the expected revenue is 2. This follows from the fact the expected revenue from each item is at most 1. If we bundle the two items together and offer the package at \$3, then the expected revenue is

$$\mathsf{Revenue} = 3 \times \mathbb{P}[v_1 + v_2 \geq 3] = 3 \times \frac{3}{4} = \frac{9}{4} > 2.$$

Hence, **bundling items may help**!

The effect of bundling becomes more obvious when the number of items becomes larger. Since they are i.i.d., by the Central Limit Theorem (or Chernoff bound), we know that the bidder's value for the grand bundle (containing everything) will essentially be a random variable drawn from a Gaussian distribution with mean $n \cdot \mu$ and variance $n \cdot \sigma^2$, where μ and σ is the mean and variance of the value for a single item.

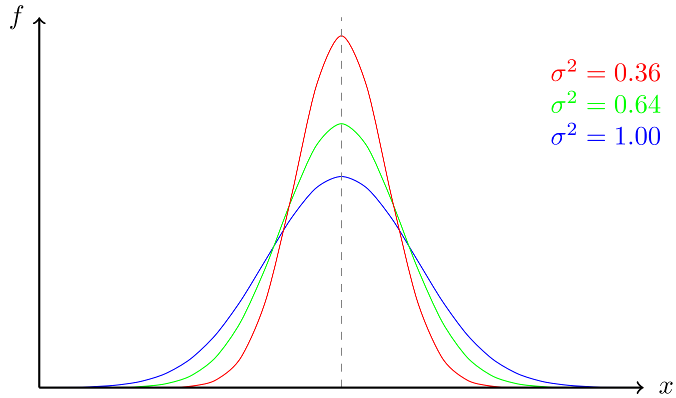

Since the ratio between the standard deviation and the mean drops quickly (in the speed of $\sqrt{n}$). If we set the price to be $(1-\epsilon)n \cdot \mu$, then the bidder will buy the grand bundle with probability almost 1. Thus, revenue is *almost* the expected value. This is the best we can hope for.

2.2 Second example

Now, suppose we take $\mathcal{F} = U\{0,1,2\}$ instead. Selling the items separately gives a revenue of \$4/3. The best way to sell the grand bundle is set it at price \$2, which gives the same revenue. But what are other ways to sell the items? Consider the option of either buying of the two items for \$2, or buying both items for \$3. In this case, the bidder's choice is given by the following table:

$v_1 \backslash v_2$	0	1	2
0	\$0	\$0	\$2
1	\$0	\$0	\$3
2	\$2	\$3	\$3

Each entry has probability 1/9 and the amount assumes the bidder plays in the best possible way. Hence, the expected revenue is

$$\text{Revenue} = 3 \times \frac{3}{9} + 2 \times \frac{2}{9} = \frac{13}{9} > \frac{4}{3}.$$

2.3 Third example

Take two distributions $\mathcal{F}_1 = U\{1,2\}$ and $\mathcal{F}_2 = U\{1,3\}$. Consider the option of either buying both items for \$4 or running a "lottery" where the bidder gets the first item for sure, gets the second item with probability 1/2 and pay \$2.50. The revenue in this case is just \$2.65.

Every deterministic auction — where every outcome awards either nothing, the first item, the second item, or both items — has strictly less expected revenue. Therefore, **randomization may help**!

3 Unit-demand Bidder Pricing Problem (UPP)

Consider the following fundamental pricing problem depicted in the figure below:

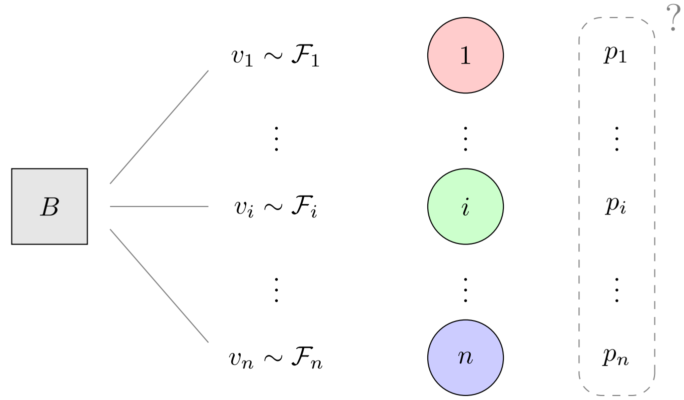

We want to pick the item (e.g. cellphone) that maximize the utility of the bidder B, who values each item i with v_i drawn from some distribution $\mathcal{F}_i$. The retailer has to come up with n prices $p_1, \ldots, p_n$. The bidder chooses the item that maximizes $v_i - p_i$, if any of them is positive. The revenue will be the corresponding p_i. We want to focus on pricing only, not considering randomized ones. In this setting, it is important to note that using randomization can only get a constant factor better than pricing.

3.1 Our goal for UPP

We want to design a pricing scheme that achieves a constant fraction of the revenue that is achievable by the optimal scheme. Hence, we want to design a polynomial time algorithm that takes the $\mathcal{F}_i$'s as input and outputs a pricing scheme $p_1, \ldots, p_n$. We shall assume that all $\mathcal{F}_i$'s are regular, that is, the virtual value function φ_i is non-decreasing in v_i.

Theorem (CHK '07)**.** *There exists a simple pricing scheme (polynomial-time computable) that achieves at least $\frac{1}{4}$ of the revenue of the optimal scheme.*

Remark. The constant can be improved with a better analysis.

3.2 Benchmark

We have seen approximation results before in Lecture 6,7 for simple nearly-optimal auctions. However, the benchmark was clear there – Myerson's auction. In the current settings however, we do not have a good grasp on what the optimal pricing scheme is or even how it looks like. We wish to compare with the optimal revenue, but we have no clue what the optimal revenue is. Do we know any natural upper bound for the optimal revenue? Again, it is not immediately obvious how to proceed.

3.3 Two scenarios

We consider two scenarios:

($\mathcal{A}$) *First scenario.*

- One unit-demand bidder B
- n items
- Bidder's value for the i-th item v_i is drawn independently from $\mathcal{F}_i$.

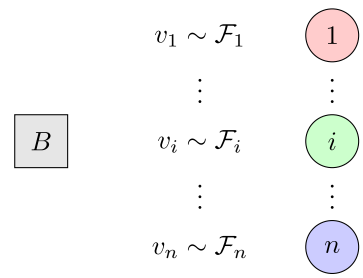

($\mathcal{B}$) *Second scenario.*

- n bidders
- One item I
- Bidder i's value v_i for the item is drawn independently from $\mathcal{F}_i$.

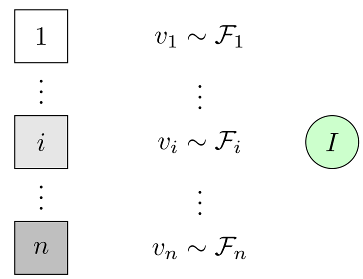

Here is a lemma that draws an important relationship between the two scenarios described above.

Lemma. *The optimal revenue achievable in scenario $\mathcal{A}$ is always less than the optimal revenue achievable in scenario $\mathcal{B}$.*

Proof: Let $\mathsf{PRev}(p)$ be the expected revenue in $\mathcal{A}$ with pricing p. Further, let $\mathsf{ARev}(M)$ be the expected revenue in $\mathcal{B}$ with auction M. For any p, we construct a mechanism M_p for $\mathcal{B}$ as follows:

We first set p_i as the reserve price for bidder i, and give the item to bidder

$$i^* = \arg\max_i{}^+ (b_i - p_i).$$

This is our allocation rule: give the item to the bidder with largest surplus $b_i - p_i$. It is clear why this is monotone. Thus, the mechanism is DSIC. The payment rule is given by

$$\mathsf{Payment} = p_{i^*} + \underbrace{\max_{j \neq i^*}(b_j - p_j)}_{\geq 0} \geq p_{i^*}.$$

If $\mathbf{v}$ is the valuation in $\mathcal{A}$, we buy item i^*, in which case the payment is just p_{i^*}. Therefore, the payment in scenario $\mathcal{A}$ is always less than the optimal revenue achievable in scenario $\mathcal{B}$, as desired. $\square$

Remark. This means the revenue of Myerson's auction in scenario $\mathcal{B}$ is an upper bound for the revenue in scenario $\mathcal{A}$, and we can use it as a benchmark.