Algorithmic Game Theory June 1, 2016

Lecture 11

Lecturer: Michal Feldman Scribe: Chen Shterental, Jonathan Yulzary, Tzahi Taub, Michael Shkolnik, Yotam Feldman, Nimrod Rappoport, Yuval Globerson

1 Introduction

Instead of focusing on the design of mechanisms that maximize the social welfare,

$$SW = \sum_{i=1}^{n} v_i x_i$$

over all feasible outcomes $(X_1, X_2, \ldots, X_n)$ in some set X, this lecture we'll focus on auctions that are explicitly designed to raise as much seller's revenue as possible.

2 Revenue maximization

Example 1 *Single Seller, Single Bidder, Single Item: The bidder has a private valuation v with distribution F, and the seller doesn't know v, but only F. If the seller knew v, he would choose the price r=v. This auction is called "posted prices", or take-it-or-leave-it offers.*

But with v private, what should we do? It's not obvious how to reason about this question. If the seller posts a price of r, then its revenue is either r (if $v \geq r$) or 0 (if $v < r$). 'If v is private, and we want to maximize the SW, we should choose r=0. If we want to maximize revenue, as we already said, this is not so easy. The fundamental issue is that, for the revenue objective, different auctions do better on different inputs.

For example, a posted price of 100 will do very well on inputs where v is 100 or a little larger, and terribly on smaller inputs (for which smaller posted prices will do better). Indeed, if the posted price is 100, and v=500, the revenue isn't maximized.

2.1 Bayesian Model

Our model comprises the following ingredients:

- single parameter environment.

- the private valuation v_i of participant i is assumed to be drawn from a distribution F_i, with density function f_i with support contained in $[0, v_{max}]$. We assume that the distributions $F_1, \dots, F_n$ are independent(but not necessarily identical). Recall that $F_i(z)$ denotes the probability that a random variable drawn from F_i, has a value at most z.
- The distributions $F_1, \dots, F_n$ are known in advance to the seller, but the realizations $v_1, \dots, v_n$ of bidder's valuations are private.
- We'll focus on DSIC auctions.

Definition 1 *"optimal" auction is the auction that, among all DSIC auctions, has the highest expected revenue, where the expectation is with respect to the given distribution $F_1 \times \dots \times F_n$ over valuation profiles v (assuming truthful bidding).*

Example 2 *Single Seller, Single Bidder, Single Item, Revisited:*

We have a private valuation v, with a distribution F, which is known to the seller, but v isn't. The expected revenue of a posted price r, is simply: $E[Rev] = r \cdot (1 - F(r))$ when r is the revenue of a sale, and $1 - F(r) = Pr(v \geq r)$, as $F(r) = Pr(v \leq r)$.

Suppose, for example, that $F = U[0, 1]$, i.e. $F(r) = r$, Then the optimal price(also called the monopoly price) is r=0.5, since it maximizes $E[Rev] = r \cdot (1 - r)$ and the expected revenue is 0.25.

Example 3 *Single-item auction, two bidders with distributions $F_1, F_2 = U[0, 1]$:*

We could, of course, run the Vickrey auction. It's revenue is the expected value of the smaller bid, which is $\frac{1}{3}$. Indeed, $E_{F1,F2=U[0,1]}[min(v_1, v_2)] = \frac{1}{3}$.

We could also supplement the Vickrey auction with a reserve price, analogous to the "opening bid" in an eBay auction. In a Vickrey auction with a reserve price r, the allocation rule awards the item to the highest bidder, unless all bids are less than r, in which case no one gets the item. The corresponding payment rule charges the winner(if any) the $max\{second - highest - bid, r\}$.

For example, r=0.5.

Let's divide to three cases: $b_1, b_2 > \frac{1}{2}, b_1, b_2 < \frac{1}{2}$, and one of $b_i > \frac{1}{2}$ and the other is smaller than $\frac{1}{2}$:

- *$b_1, b_2 > \frac{1}{2}$, the probability for this event is 0.25, and $E_{F1,F2=U[\frac{1}{2},1]}[min(v_1, v_2)] = \frac{2}{3}$.*
- *$b_1, b_2 < \frac{1}{2}$, the probability for this event is 0.25, and the revenue is 0.*
- *One of $b_i > \frac{1}{2}$ and the other is smaller than $\frac{1}{2}$, the probability for this event is 0.5, and the revenue is $max(< 0.5, 0.5) = \frac{1}{2}$.*

Therefore, $E[Rev] = 0.25 * \frac{2}{3} + 0.25 * 0 + 0.5 * 0.5 = \frac{5}{12} > \frac{1}{3}$.

Fortunately, We succeeded in increasing the expected revenue, from that of a Vickrey auction. We'll learn later that this is the auction which maximizes revenue,among all DSIC auctions, of that kind.

2.2 The Revelation Principle

Recall the definition of DSIC mechanisms:

Definition 2 *A mechanism is DSIC if the following conditions hold:*

- *Each player has a dominant strategy.*
- *It's an equilibrium when the strategy of each player is direct revelation.*

If we remove condition 1, i.e., there exists an equilibrium, when each player tells his private value, but it isn't a dominant strategy, then the size of our mechanisms space grows.

If we remove condition 2, then the size of our mechanisms space doesn't grow.

Formally:

Theorem 3 *Let M be a mechanism, for which each player has a dominant strategy. Then, there exists an equivalent mechanism M'(a mechanism which yields the same result), for which the dominant strategy of each player is to reveal his private value.*

Proof:

Given a mechanism M, M' will simulate M on $s_1(v_1), \ldots, s_n(v_n)$, when $s_i(v_i)$ is the dominant strategy of player i. M' is DSIC: Since no player i can increase his profit by bidding $b_i \neq s_i(v_i)$, then also in M', no bidder i can increase his profit by bidding $b_i \neq v_i$. See figure 1.

■

Corollary 4 *If M is a mechanism, for which each player has a dominant strategy, then we can assume w.l.o.g that M is DSIC.*

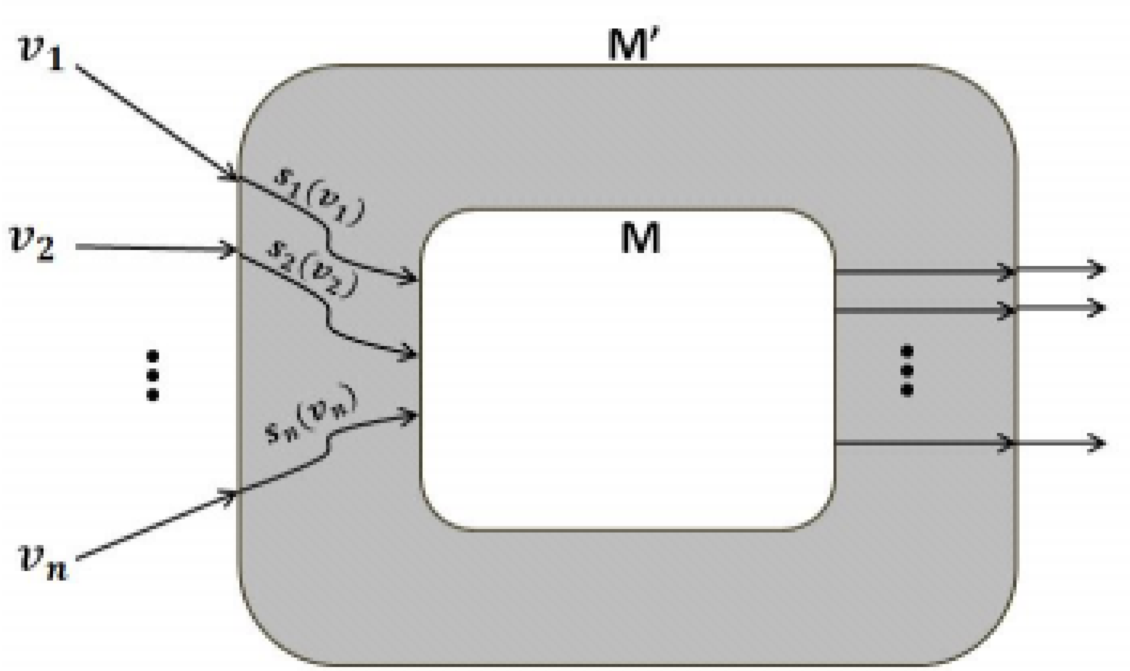

Figure 1: The design of M'

3 Looking for the optimal auction

F_i, f_i, v_i etc. as used before. $x_i(b_i, b_{-i})$ denotes the allocation to player i with bids b. We will look for an auction which maximizes the expected revenue, which is

$$\mathbb{E}_{v\sim F}[\sum_{i=1}^{n} p_i(v)] \tag{1}$$

where F = $F_1 \times F_2 \times ... \times F_n$.

Using Myerson's equation: $p_i(b_i, b_{-i}) = \int_0^{b_i} z x_i'(z, b_{-i})dz$. Now, let us fix i, v_{-i}.

$$\mathbb{E}_{v_i\sim F_i}[p_i(v)] = \int_0^{v_{max}} p_i(v) f_i(v_i) dv_i = \int_0^{v_{max}} (\int_0^{v_i} z x_i'(z, v_{-i})dz) f_i(v_i) dv_i$$

changing integration order leads to

$$= \int_0^{v_{max}} (\int_z^{v_{max}} f_i(v_i) dv_i) z x_i'(z, v_{-i}) dz$$

using the definition of a CDF function and then integration by parts:

$$= \int_0^{V_{max}} (1 - F_i(z)) z x_i'(z, v_{-i}) dz = \overbrace{[(1 - F_i(z)) z x_i(z, v_{-i})]\Big|_0^{v_{max}}}^{0} - \int_0^{v_{max}} x_i(z, v_{-i})[1 - F_i(z) - z f_i(z)]dz =$$

$$\int_0^{v_{max}} \underbrace{(z - \frac{1-F_i(z)}{f_i(z)})}_{\Phi_i(z)} x_i(z, v_{-i}) f_i(z) dz$$

If we would have replaced the part denoted by $\Phi_i(z)$ with z we would get the *expected social welfare of player i*, hence we will call Φ the *Virtual Welfare*.

Example 4 *For* $F_i = Uniform(0,1)$ *we get* $\Phi_i(v_i) = v_i - \frac{1-v_i}{1} = 2v_i - 1$

Assigning some numbers, we get for example: $\Phi_i(1) = 1$, $\Phi_i\left(\frac{1}{2}\right) = 0$, $\Phi(0) = -1$

*Notice that **virtual welfare may be negative**.*

From the above we get the expected payment of player i:

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(v)\right] = \mathbb{E}_{v_i \sim F_i}\left[\Phi_i(v_i) \cdot x_i(v)\right]$$

Now, take expectation over the values of the rest of the players:

$$\mathbb{E}_{v \sim F}\left[p_i(v)\right] = \mathbb{E}_{v \sim F}\left[\Phi_i(v_i) \cdot x_i(v)\right]$$

Finally, sum over all players:

$$\mathbb{E}_v\left[\sum_{i=1}^n p_i(v)\right] = \sum_{i=1}^n \mathbb{E}_v\left[p_i(v)\right] = \sum_{i=1}^n \mathbb{E}_v\left[\Phi_i(v_i) \cdot x_i(v)\right] = \mathbb{E}_v\left[\sum_{i=1}^n \Phi_i(v_i) \cdot x_i(v)\right]$$

When the first and last equalities are by the linearity of expectation.

From the two ends of the above equation, we get by definition:

Corollary 5

$$\mathbb{E}\left[Revenue\right] = \mathbb{E}\left[Virtual\, welfare\right]$$

Example 5 *An auction where there is a single product and n players whose value distributes i.i.d according to a distribution function F. In order to maximize the expected revenue we need to maximize the expected virtual welfare:*

$$\mathbb{E}_{v \sim F}\left[\sum_{i=1}^n \phi_i(v_i) x_i(v)\right] \tag{2}$$

Given v, we look for x(v) that maximizes $\sum_{i=1}^n \phi_i(v_i) x_i(v)$, *under the constraint that* $x_i(v) \in \{0,1\}$ *and* $\sum_{i=1}^n x_i(v) \leq 1$.
In the optimal allocation function, x(v), the product will therefore be given to the player with highest $\phi_i(v_i)$, *if* $\phi_i(v_i) \geq 0$. *Otherwise, the product won't be sold. Note that if the virtual welfare is a monotone function, this allocation function is also monotone.*

Definition 6 *A distribution F is regular if the virtual value, $v - \frac{1-F(v)}{f(v)}$, is monotonically increasing in v.*

Corollary 7 *For a regular distribution, the above allocation function is monotone. Therefore, according to Myerson,there is a payment rule for that allocation function such that the mechanism is DSIC.*

The optimal mechanism is therefore a second price auction with a minimal price which equals $\phi^{-1}(0)$. For example, if $F = U[0, 1]$ (F distributes uniformly between 0 and 1), the optimal mechanism is a second price auction with a minimal price which equals $\frac{1}{2}$.

The above allocation depends on the value distribution F. However, if the players have different value distributions (even though they're all regular and independent!), the allocation function may be "weird". For example, we may get an allocation function in which the highest bidder doesn't necessarily win the product. The optimal allocation is not necessarily a "simple" one. Who wants to participate in such a strange auction? We will therefore look for "simple", DSIC mechanisms, where the expected revenue will approximate the optimal expected revenue.

4 Prophet Inequality

In this section we consider a game with n steps. In each stage i the player is offered a prize $\pi_i \sim G_i$. The player must then choose whether to accept the prize and end the game, or to continue to the next levels.

The distributions $G_1, \ldots, G_n$ are known in advance to the player and are independent. The player sees the realization of π_i only on step i.

This game is a form of *stopping problem* which are also common in life (choosing life-partners, *Who Wants to Be a Millionaire?*, ...).

We can solve this game using backward induction: when the last step of the game is reached it is of course always beneficial to take the suggested prize (because otherwise we don't win anything). In step $n-1$ we should take the suggested π_{n-1} iff it is larger than the expected prize in the future steps, in this case just step n. We can continue this reasoning to decide whether to take π_{n-2} since we have already decided the strategy for steps $n-1$ and n, so we can compute the expected prize from continuing in the game, and so on.

This strategy is relatively hard to compute and hard to describe, and is also sensitive to small changes in the distributions $G_1, \ldots, G_n$. We would like to have "simple" strategy that is also quite good, in the sense that although it may not be optimal it's still an approximation to the optimal strategy.

Theorem 8 *For every independent distributions $G_1, \ldots, G_n$ there exists a "simple" strategy that in expectation wins a prize of at least*

$$\frac{1}{2}\mathbb{E}_\pi\left[\max \pi_i\right]$$

By "simple" we mean that this is a threshold strategy*: it can described using a single parameter t such that in every step i we choose to accept the prize iff $\pi_i \geq t$. We will also require that computing t should be relatively simple.*

Note that $\mathbb{E}_\pi\left[\max \pi_i\right]$ is undoubtedly the highest expected value a player can get, and it's even not all clear the complicated strategy described above actually attains this bound. Still we can approximate it to a factor of 2 using a simple strategy, as follows.

Proof: Throughout this proof, for every $z \in \mathbb{R}$ denote $z^+ = \max\{z, 0\}$.

For every threshold t let $q(t)$ be the probability that the player doesn't win any prize during the game. Note that this may happen when using a simple threshold strategy if during all steps of the game $\pi_i < t$ and also in the last step $\pi_n < t$ and the player doesn't take it (although there is no reason not to — it's clear that a simple threshold strategy can't be optimal).

$q(t)$ is a monotonically increasing function of t. However, as t increases, the expected prize that the player wins increases, conditioned on the event that the player actually wins a prize. The question is how to choose a good t for this trade-off.

We would like to bound the expected prize for the player. The expected prize is at least

$$q(t) \cdot 0 + (1 - q(t)) \cdot t$$

because with probability $1 - q(t)$ the player wins a prize, and according to the threshold strategy it must be at least t.

This bound is very loose because when the player wins a prize it may be greater than t. If that there is a single step i in which $\pi_i \geq t$ the prize will be π_i, not just t, so we need to add $\pi_i - t$ to the bound. If there is more than one such step we must add the first of them and the reasoning becomes more complicated, so we will consider the bonus only in cases where there is just a single step in which the prize was at least t; if there's more than one we use the simple bound of t.

Therefore, the expected prize by playing the threshold strategy with threshold t is at least

$$(1 - q(t)) \cdot t + \sum_{i=1}^{n} \mathbb{E}\left[\pi_i - t \mid \pi_i \geq t,\ \forall j \neq i.\ \pi_j < t\right] \cdot \Pr\left[\pi_i \geq t\right] \cdot \Pr\left[\pi_j < t\ \forall j \neq i\right].$$

To be continued in the next lecture. ■