Lecture 12

Lecturer: Michal Feldman Scribe: Tsvi Fominikh, Oren Gilon, Guy Shalev, Ilia Gorelik, Slava Novgorodov

1 Bayesian Truthfulness

In this section, we will discuss games in the Bayesian setting and define a new concept of truthfulness for these games. We will then find auctions that maximize the expected revenue.

Definition 1 *A mechanism X, P is called Bayesian-Truthful if the profile where each player tells the truth is a Bayesian Nash Equilibrium (BNE) in the game induced by the mechanism.*

It is important to differentiate between the two concepts of truthfulness we have seen. Throughout this course, we have looked at Dominant-Strategy-Truthfulness, where telling the truth is a dominant strategy for each player. We are now looking at a weaker sense of truthfulness, where telling the truth is simply a best-response to all other players telling the truth.

Observation 2 *The Revelation Principle exists in the Bayesian setting.*

Proof: For any mechanism M where there exists a BNE $S_1(v_1), ..., S_n(v_n)$ there exists an equivalent mechanism M' that is direct-revelation i.e. a mechanism that is Bayesian-Truthful.This is easy to see through a simulation argument. We will simply create a mechanism M' that upon inputs $b_1, ...b_n$, simulates M on $S_1(b_1), ..., S_n(b_n)$. Since $S_1(v_1), ..., S_n(v_n)$ is a BNE in M, $v_1, ...v_n$ is clearly a BNE for M'. ■

Observation 3 *Myerson's Lemma holds in the setting of Bayesian-Truthfulness:*

Myerson's lemma in the single-parameter setting implies that an allocation rule X is implementable $\Leftrightarrow$ X is monotone. In addition, if X is monotone, there exists a single payment function P such that (X,P) is Bayesian-Truthful.

Corollary 4 *Revenue Equivalence Theorem: Mechanisms that have the same allocation in a BNE have the same revenue.*

Proof: The revenue is a function of the payments, and the payments are only a function of the allocation, through Myerson's Lemma. ■

Example 1 *First Price Auction BNE*

Let us look at two players whose values $v_i \sim U[0,1]$. In the previous lecture, we had to guess a strategy that is a BNE. We will try to use the Revenue Equivalence Theorem to find that strategy. We will only guess that the player with the highest value will win the object. The expected payment of player i is the same as that in a second-price auction. The expected payment in a second price auction is

$$\begin{aligned} E[P_1(v_1)] &= E[P_1(v_1)|1\ wins] * Pr(1\ wins) + E[P_1(v_1)|1\ loses] * Pr(1\ loses) \\ &= E[v_2|v_2 < v_1] * v_1 \\ &= \frac{v_1}{2} * v_1 = \frac{v_1^2}{2} \end{aligned}$$

Thus we get that in a first price auction

$$\begin{aligned} v_1^2/2 = E[P_1(V_1)] &= E[P_1(v_1)|1\ wins] * Pr(1\ wins) \\ &= E[P_1(v_1)|1\ wins] * v_1 = S_1(v_1) * v_1 \end{aligned}$$

and this gives us that

$$S_1(v_1) = \frac{v_1}{2}$$

Which is exactly the strategy we found in the previous lecture. We just need to verify that our origisnal assumptions about the strategy hold. Our only assumption was that the player with the highest value has the highest bid. Indeed, if each player bids $v_1/2$, the player with the highest value wins the auction. This means the strategy we found, $S_i(v_i) = v_i/2$ is indeed a BNE. As we found the strategy through the revenue equivalence theorem, no player can improve his expected revenue, as he would then also improve his expected revenue in the BNE of the second price auction.

2 Approximate Revenue Maximization

In this section we will try to find simple auctions that approximate the maximal revenue for the auctioneer.

Reminder

We look for an auction that maximizes the expected revenue.

Theorem 5

$$E[revenue] = E[\sum_{i=1}^{n} \phi_i(v_i) * X_i(v)]$$

Where $\phi_i(v_i)$ is the Virtual Value, and is definded to be

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

Definition 6 *A distribution F is called regular if ϕ_i is monotone.*

We study a more complex settings, where the players' values are drawn from different distributions. We assume n players, where $v_i \sim F_i$, and $\{F_i\}$ are independent and regular. What does an optimal auction look like in this case? We will see two things:

1. In the case of two players, the optimal auction is not intuitive.
2. Simple auctions can provide a good approximation to the optimal revenue.

Example 2 *We consider an auction where two players bid for a single item. Their values for the item are distributed $v_1 \sim U[0,1]$, $v_2 \sim U[0,2]$. What is the optimal auction in this setting in terms of maximizing the revenue?*

We calculate the virtual values, as defined in the previous lesson. $\Phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$, where f_i is the probability density function, and F_i is the cumulative distribution function. For our $v1, v2$ we get:

$$\Phi_1(v_1) = v_1 - \frac{1 - v_1}{1} = 2v_1 - 1$$
$$\Phi_2(v_2) = v_2 - \frac{1 - \frac{v_2}{2}}{\frac{1}{2}} = 2v_2 - 2$$

When is it optimal to give the item to player 1? When:

$$\Phi_1(v_1) > \max(\Phi_2(v_2), 0) \iff v_1 > v_2 - \max(v_2 - \frac{1}{2}, \frac{1}{2})$$

Similarly, it is optimal to give the item to player 2 when:

$$\Phi_2(v_2) > \max(\Phi_1(v_1), 0) \iff v_2 > v_1 - \max(v_1 + \frac{1}{2}, 1)$$

So the allocation function can be expressed by the following graph, where the axes of the graph represent the values of the random variables v_1, v_2:

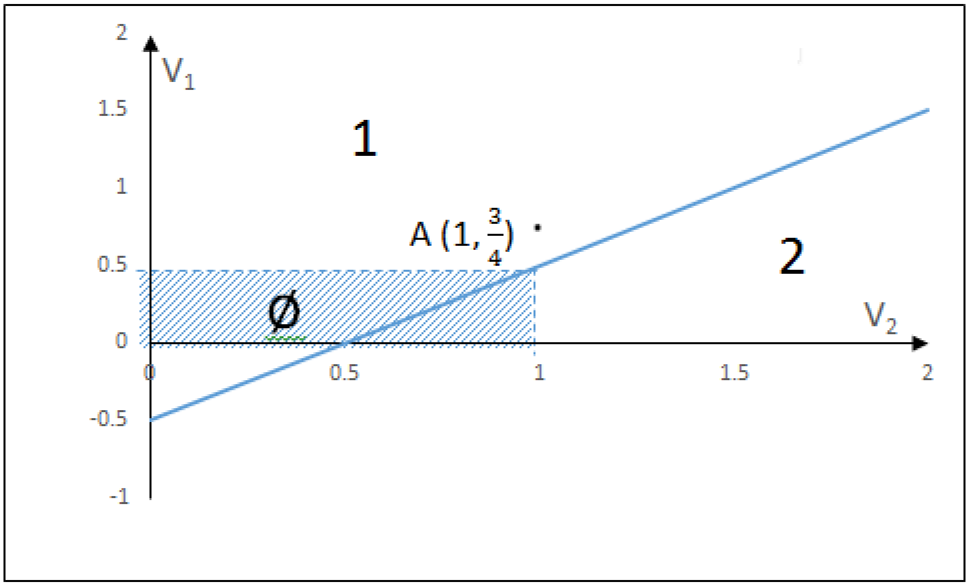

The allocation function descibed using a graph

As can be seen, the optimal solution for this simple auction is not simple at all, and does not have an elegant representation. Also, as can be seen in the graph at point A, the item is given to player 1 even though he values it less than player 2. It is interesting to investigate simpler, non-optimal auctions, that are a good approximation to the optimal auction. Previously, we used approximized algorithms to overcome computational complexity, but now we use them to overcome the "complexity of the auction rules".

3 Distribution-based Mechanism Design

In this section, we will cover the following scenarios regarding the distribution which the players' valuation are sampled from:

- Bayesian Approximation - The distribution is assumed to be known by both the designer of the mechanism, and the analyser of the game.
- Prior Independent Approximation - The distribution is unknown to the designer, but is known to the analyser.
- Prior Free Approximation - The distribution is unknown to both the designer and the analyser.

3.1 Bayesian Approximation

We will now see a very elegant result from the world of optimal stopping.

Problem 1 *A player is playing a game that lasts n days. Every day he is offered an item, with a value distibuted* $0 \leq \Pi_i \sim F_i$, $1 \leq i \leq n$. *On the i'th day, the realization of* Π_i *is revealed. The player may accept the item and end the game, or reject it and continue to the next day. Once rejected, an offer cannot be recalled. What is the optimal strategy, in terms of the expected value for the player?*

Solution 1 *The optimal solution can be achieved by the backward-induction method. If the player has reached the last day, the optimal behavtiour is trivial - he takes the item and the value expectation is* $\mathbb{E}[F_n]$. *The day before, the player sees the realization of* Π_{i-1}. *He should accept the item iff* $\Pi_{i-1} \geq \mathbb{E}[F_n]$. *According to this dominating strategy,* $\mathbb{E}[F_{n-1}]$ *can be calculated, and so on by induction.*

The strategy mentioned is in fact the optimal strategy. The problem is it is quite complicated and is not resistible to minor changes in the problem, such as changes to some of the distributions F_i, or a permutation on the distributions. We will now see a strategy, that is simpler, more change-resistable and still efficient.

Definition 7 *For any independent distributions* $F_1, ..., F_n$, *a "t-threshold" strategy is defined by the rule "accecpt item* $\Pi_i \iff \Pi_i \geq t$*"*

Theorem 8 *For any independent distributions* $F_1, ..., F_n$, *there exists some* $t \in \mathbb{R}$ *so that the t-threshold strategy achieves a value expectation of at least* $\frac{1}{2}\mathbb{E}_\Pi[\max_{1 \leq i \leq n} \Pi_i]$

This result is known as the prophet-inequality, because the value expectation achieved is at least half the value that would be achieved by a "prophet-player", that knows all the realizations in advance.

Proof: For simplicity we define $z^+ = max\{z, 0\}$.
Let $p(t) = Prob[\Pi_i < t|v_i]$ be the propability of not getting any item at all.
Then with probability $p(t)$ the profit will be zero and with probability $[1 - p(t)]$ the profit will be at least t.
In addition we notice that if exectly one item π_i has a value greater then t the profit will be Π_i increasing the profit by $\Pi_i - t$. Summarizing those three cases we obtain profit of at least:

$$\mathbb{E}[\text{Payoff}_t] \geq$$

$$0 \cdot p(t) \; + \; t \cdot [1 - p(t)] \; + \; \sum_{i=1}^{n} \mathbb{E}[\Pi_i - t | \Pi_i \geq t, \forall j \neq i. \Pi_j < t] \cdot Pr[\Pi_i \geq t] \cdot Pr[\Pi_j < t | \forall j \neq i] =$$

$$t \cdot [1 - p(t)] \;+\; \sum_{i=1}^{n} (\mathbb{E}[\Pi_i - t | \Pi_i \geq t] \cdot Pr[\Pi_i \geq t]) \cdot Pr[\Pi_j < t | \forall j \neq i] =$$

$$t \cdot [1 - p(t)] \;+\; \sum_{i=1}^{n} \mathbb{E}[(\Pi_i - t)^+] \cdot \frac{p(t)}{Pr[\Pi_i < t]} \geq$$

$$t \cdot [1 - p(t)] \;+\; \sum_{i=1}^{n} \mathbb{E}[(\Pi_i - t)^+] \cdot p(t)$$

On the other hand

$$\mathbb{E}[\max_{1 \leq i \leq n} \Pi_i] = \mathbb{E}[t + \max_{1 \leq i \leq n} (\Pi_i - t)] \leq$$

$$t + \mathbb{E}[_{1 \leq i \leq n} (\Pi_i - t)^+] \leq$$

$$t + \sum_{i=1}^{n} \mathbb{E}[(\Pi_i - t)^+]$$

By choosing t so that $p(t) = \frac{1}{2}$ we obtain:

$$\mathbb{E}[\text{Payoff}_t] \geq t \cdot [1 - \frac{1}{2}] \;+\; \sum_{i=1}^{n} \mathbb{E}[(\Pi_i - t)^+] \cdot \frac{1}{2} \geq \frac{1}{2} \mathbb{E}[\max_{1 \leq i \leq n} \Pi_i]$$

■

Remark: This proof doesn't shows a tight bound since in all the cases where more then one item has value greater then t we assume a profit of exactly t and not the real value.

Problem 2 *We now return to our single-item auction with n bidders and valuations drawn from (not necessarily identical) regular distributions $F_1, \ldots, F_n$. We regard the virtual valuation $\phi_i(v_i)^+$ of a bidder as ith prize in the n days game.*

Recall that the Revenue of the optimal auction is

$$\mathbb{E}_v[\sum_{i=1}^{n} \phi_i(v_i) \cdot X_i(v)] = \mathbb{E}[\max_i \phi_i(v_i)^+]$$

Our goal is to design a simple auction with a revenue close to the optimal.

Solution 2 *We think about our action as an n days game with $G_i = F_i$. From the prophet-inequality we know that there exists an "t-threshold" strategy that will give at a revenue of at least $\frac{1}{2}$ of the optimal auction.*

1. *We will choose t so that $Pr[\max_i \phi_i(v_i) \geq t] = \frac{1}{2}$*
2. *We will give the item to player i if $\phi_i(v_i) \geq t$ (if such exists). Ties will be broken arbitrary as long as the monotonicity is satisfied.*

Prophet Inequality implies that every auction with the same allocation rule satisfies

$$\mathbb{E}_v[\sum_{i=1}^{n} \phi_i(v_i) \cdot X_i(v)] \geq \frac{1}{2}\mathbb{E}_v[\max_{i=1}^{n} \phi_i(v_i) \cdot X_i(v)]$$

This proves that the $t-threshold$ auction gives a revenue of at most $\frac{1}{2}$ of the optimal auction

Example 3 *Here is a specific such allocation rule:*

For each player we will define minimum price $v_i = \phi_i^{-1}(t)$ (with the same t from above). We will give the item to the highest bidder that meets his minimum price (v_i). Then:

From Problem 2 we know that this auction's revenue is at least $\frac{1}{2}$ of the optimal auction's revenue, but this is a much more simple auction.

3.2 Prior-Independent Approximation

We now discuss Prior-Independent Approximation, defined earlier.

Reminder

Definition 9 *Prior Independent Approximation - The distribution is unknown to the designer, but is known to the analyser.*

This can be relevant in a narrow market, where it is difficult to obtain statistics about the distributions of buyers. We look at two situations:

Situation A n players, having valuations which are i.i.d. from distribution F, with an optimal auction. In this case, we note that the optimal auction strongly depends on F.

Situation B $n + 1$ players, having valuations which are i.i.d. from distribution F, with a second-price auction. In this case, the auction does not depend on F.

We will show that if we can add another player with distribution F to the simple second-price auction, we get a revenue not less than the revenue of the optimal auction. According to the situations, we define two auctions:

Auction A An optimal auction on n players having valuations which are i.i.d. from distribution F.

Auction B A second-price auction on $n+1$ players having valuations which are i.i.d. from distribution F.

Theorem 10 *If F is regular, then $E[Rev(B)] \geq E[Rev(A)]$*

Proof: For the proof, we define another auction C:

Auction C An auction with $n+1$ players. The auction simulates the optimal auction on the first n players. If no player got the product in the simulated auction, the auction gives the product to player $n+1$ for free.

By definition of C, it follows that $E[Rev(C)] = E[Rev(A)]$.
The auction C that we defined is an auction with $n+1$ players, in which there is always a player who gets the product. Therefore, the revenue of Auction C is at most the revenue of the optimal auction on $n+1$ players where one of the players always gets the product. We try to formulate what this optimal auction is:
We need an auction which gives the product to a player with the highest virtual value $\phi_i(v_i)$. F is regular, so $\phi_i(v_i)$ is non-decreasing in v_i. Therefore, the player with the highest $\phi_i(v_i)$ is actually the player with the highest v_i. The auction we looked for is exactly a second-price auction (Auction B).
We get the required inequality:

$$\overbrace{E[Rev(B)]}^{\substack{\text{An optimal auction with n+1 players} \\ \text{which always gives away the product}}} \geq \overbrace{E[Rev(C)]}^{\substack{\text{Some auction with n+1 players} \\ \text{which always gives away the product}}} = E[Rev(A)]$$

■

3.3 Prior-Free Approximation

We will now discuss Prior-Free Approximation.

Reminder

Definition 11 *Prior Free Approximation - The distribution is unknown to both the designer and the analyser.*

First of all, when discussing good approximations in this case, we have to define what "good" actually means.

Reminder

In the previous cases, we looked for a good approximation to the optimal auction.

In the current setting, we cannot assume that a distribution for the bidder valuations even exists. Thus, we don't have a general optimal auction, so the criteria from the previous cases is irrelevant. Instead of that, we define the following requirement from a good approximation:

Requirement A good approximation in a Prior-Free setting will also give us a good approximation in a Prior-Independent setting. That means, it will be a good approximation to the optimal auction, for every distribution.

How can we tell if an approximation meets this requirement? We will be looking for an appropriate *benchmark* - an auction to compare our approximation mechanisms to. We require the following property: a mechanism which will give a good approximation to the revenue of our benchmark is promised to give a good approximation to the revenue of the optimal Bayesian mechanism, for every distribution. We will define a specific type of auctions and show that such auction can serve as a benchmark.

Definition 12 *An auction is envy-free if no player prefers the allocation and payment of another player over his own allocation and payment.*

Example 4 *We look at digital goods, which are goods with no production costs.*

Assume that we have n people with values $v_1, ..., v_n$, with all $v_i \geq 0$, and one product with no production costs. Then, it is easy to maximize social welfare - every player will get the product for free.
What about maximizing revenue? We take the values in a sorted order, $v_{(1)} \geq ... \geq v_{(n)}$, and try to design an envy-free auction for this case. First assume that all players have the same distribution F. Then our auction is:

Envy-free auction 1 We offer each player a take-it or leave-it deal, with price $\phi^{-1}(0)$.

This is an envy-free auction, as all players are in the same situation and no player will prefer to switch. Generally, if we offer all players the same price, the auction will be envy-free.

Envy-free auction 1 would work if we had information about F, but it is not true in the prior-free case which we discuss.

Observation 13 *Given a vector of valuations $v = (v_1, ..., v_n)$, the maximal revenue that can be obtained from an identical price is $\max_i i * v_{(i)}$.*

The observation follows from the fact that exactly i players will agree to an offer $v_{(i)}$ (according to the sorted order).

Envy-free auction 2 We offer each player a take-it or leave-it deal, with price v_j where $j = \arg\max_i i * v_{(i)}$.

Every mechanism which gives a good approximation to the revenue of *Envy-free auction 2* (which is $\max_i i * v_{(i)}$), will give a good approximation to the optimal Bayesian mechanism for every distribution.

Corollary 14 *Envy-free auction 2 meets our requirement and can serve as a benchmark.*

In the next lesson, we will obtain the following results:

Good result In a case where there are always at least two winners in the auction, there is an auction mechanism which provides a good approximation.

Bad result If there is a single player with a very high value, it is impossible to make him reveal this value.