CS880: Algorithmic Mechanism Design

Lecture 4: Simplicity vs Optimality **Scribe:** Ashwin Maran

4.1 Recap

In the last lecture, we studied the problem of maximizing revenue in the Bayesian setting. Given n buyers whose valuations v_i are drawn independently from known distributions F_i, we wanted to find the BIC mechanism that maximizes the auctioneer's revenue. We proved the following theorems.

Theorem 4.1.1 *Expected revenue of any BIC mechanism is equal to the expected virtual surplus.*

$$Rev(x) = \mathbb{E}_{v \sim F_1 \times \cdots \times F_n} \left[\sum_{i \in [n]} x_i(v_i) \phi_i(v_i) \right]$$

where the virtual surplus $\phi_i(v_i)$ is defined as

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

Theorem 4.1.2 *When all buyers have regular value distributions (i.e. $\phi_i(\cdot)$ is a monotone non-decreasing function), pointwise maximization of virtual surplus gives a BIC mechanism.*

This showed us that when the buyers have regular value distributions, the optimal mechanism is the mechanism that pointwise maximizes the virtual surplus. In the non-regular case, we looked at the process of ironing the revenue curve, and the ironed virtual surplus $\overline{\phi_i}$. We then studied Myerson's mechanism.

4.1.1 Myerson's Mechanism

Consider the following mechanism.

Myerson's mechanism:

1. Pointwise maximize the ironed virutal surplus.
2. In case of any ties in the ironed virtual surplus, break ties consistently.

We proved that Myerson's mechanism is the optimal revenue maximizing mechanism. We also looked at various examples, where the mechanism turned out to be very complicated, motivating the need for simpler mechanisms.

4.1.2 Vickrey auction with monopoly reserve price

We then looked at a variant of the simple Vickrey auction, where each agent has a reserve price r_i, and only agents whose value $v_i \geq r_i$ are allowed to participate. Consider the following mechanism.

Vickrey auction with monopoly reserve price:

1. Compute $r_i^* = \phi_i^{-1}(0)$.
2. Reject all agents with $v_i \le r_i^*$.
3. Run Vickrey auction over the values reported.
4. Charge critical prices $p_i = \max\left(r_i^*, \left(\max_{j \neq i : v_j \ge r_j^*} v_j\right)\right)$.

We proved the following.

Theorem 4.1.3 *In the non-iid regular setting, the Vickrey auction with monopoly reserve price gets a 2-approximation to the optimal expected revenue.*

In the rest of this lecture, we will continue to study other such auctions in the single item, regular setting, where we trade optimality for simplicity.

4.2 Posted Price Mechanism

Assume that we have n buyers each with a valuation v_i of the item, drawn independently from a known distribution F_i. In a posted price mechanism, given $F_1, \ldots, F_n$, the prices $p_1, \ldots, p_n$ for each agent is determined. Then, the buyers arrive one at a time, and the first agent whose valuation v_i exceeds the posted price p_i gets the item for the posted price p_i.

As we can see, the posted price mechanism is fairly simple. However, note that this mechanism is not making use of the competition between the users. The first agent with $v_i \ge p_i$ gets the item irrespective of the values of the remaining agents. This appears to be a serious flaw with posted price mechanisms, but as we shall see, they provide very good approximations to the optimal revenue. This is explained by the fact that although there is no competition between the users, the posted price p_i can still affected by F_{-i}.

The performance of the posted price mechanism clearly depends on the order in which the agents arrive. In a **sequential posted price mechanism**, the auctioneer is allowed to invite the agents in any specific ordering of her choosing. So, we may assume that the auctioneer invites the agents in decreasing order of their posted price. On the other hand, in an **order-oblivious posted price mechanism**, the auctioneer has no control over the order in which the agents arrive. So, in the worst case, we have to assume that the agents arrive in increasing order of their posted price. In a **random order posted price mechanism**, we assume that the agents arrive in a uniformly random order, and the auctioneer has to maximize the expected revenue.

4.2.1 Ex-ante Relaxation

The supply constraint for any single item auction is that we can allocate the item to at most one agent for any value vector $(v_1, \ldots, v_n)$. This is an ex-post feasibility constraint. The ex-ante relaxation is to replace this constraint with the following constraint.

Ex-ante relaxation: Allocate at most one item in expectation over all value vectors.

Clearly, a mechanism that satisfies the ex-ante relaxation is not necessarily feasible, since we might end up allocation more than one item to some particular value vector. However, the problem of maximizing the revenue subject to the ex-ante relaxation becomes much easier. Consider the following problem.

$$\begin{aligned} \max \quad & \sum_{i \in [n]} Z_i(q_i) \\ \text{s.t.} \quad & \sum_{i \in [n]} q_i \leq 1 \end{aligned}$$

where $Z_i(q_i)$ is defined to be the maximum revenue we can get from agent i by selling to them with ex-ante probability at most q_i. We will denote the solution to the problem by EA-OPT. Then, it is easy to see that

Lemma 4.2.1 *The optimal revenue is at most EA-OPT.*

Let us now try to solve for EA-OPT. In order to do that, we will optimize for each $Z_i(q_i)$. First let us consider the case where F_i is regular. In that case, if we want to maximize the revenue from selling with an ex-ante probability of q_i, then we want to sell the item to the agent when their value is at the highest q_i quantiles. Notice that this corresponds exactly to posting a price of $F_i^{-1}(1-q_i)$ for the agent i. In other words, when F_i is regular, $Z_i(q_i) = R_i(q_i)$.

In the non-regular setting, it becomes a little more complicated.

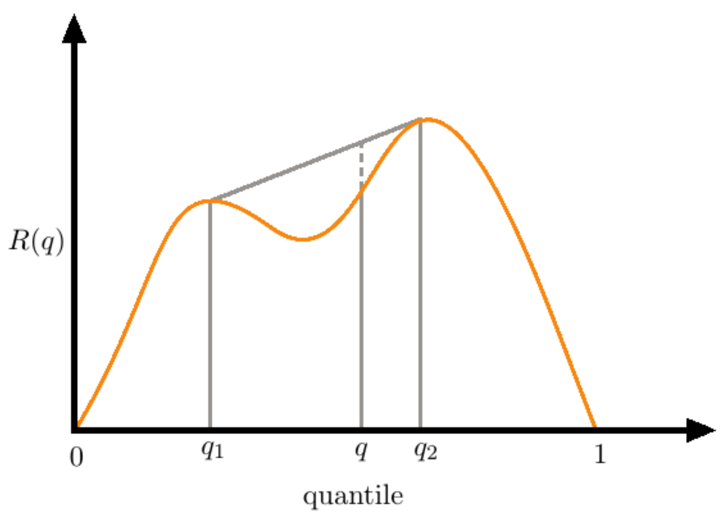

Figure 4.2.1: The revenue curve of an irregular distribution

Notice that for an ex-ante probability of q, instead of setting a posted price of $F_i^{-1}(1-q)$, the auctioneer is better off using a randomized auction. If $q = \alpha q_1 + (1-\alpha)q_2$, then the auctioneer could post a price of $F_i^{-1}(1-q_1)$ with probability α, and a price of $F_i^{-1}(1-q_2)$ with probability $1-\alpha$. Therefore, in the non-regular setting, $Z_i(q_i) = \overline{R_i}(q_i)$, and in the regions where $R_i(q_i) \neq \overline{R_i}(q_i)$, the ex-ante optimal mechanism randomizes between two posted prices.

Consider the following order-oblivious posted price mechanism.

OPM Mechanism using Ex-ante relaxation:

1. Suppose $q_1^*, q_2^*, \ldots, q_n^*$ is the optimal solution to EA-OPT.
2. Set $q_i = \frac{q_1^*}{2}$ and $p_i = F_1^{-1}(1-q_i)$.
3. Offer prices p_i to the agents in any arbitrary sequence.

Theorem 4.2.2 *In a single item auction with n agents each having non-iid valuations v_i drawn from distributions F_i, the above order-oblivious posted price mechanism gets a 4-approximation to the optimal mechanism.*

Proof: Note that

$$\begin{aligned}
\text{EA-OPT} &= \sum_{i\in[n]} Z_i(q_i^*) \\
&= \sum_{i\in[n]} \overline{R_i}(q_i^*) \\
&\leq 2\sum_{i\in[n]} \overline{R_i}(q_i) \\
&= 2\sum_{i\in[n]} Z_i(q_i)
\end{aligned}$$

The inequality is explained by the fact that $\overline{R_i}(q_i)$ is a concave function, which implies that

$$\overline{R_i}(q_i) = \overline{R_i}\left(\frac{(q_i^* + 0)}{2}\right) \geq \frac{1}{2}\left(\overline{R_i}(q_i^*) + \overline{R_i}(0)\right) = \frac{1}{2}\overline{R_i}(q_i^*)$$

Moreover,

$$\sum_{i\in[n]} q_i = \frac{1}{2}\sum_{i\in[n]} q_i^* \leq \frac{1}{2}$$

So, the total probability of selling the item is less than $\frac{1}{2}$. Therefore, for any agent, the probability that the item will still be available when they arrive is at least $\frac{1}{2}$. Therefore,

$$
\begin{aligned}
\text{ALG} &= \sum_{i\in[n]} \Pr\left[\text{item is offered to agent } i\right] \cdot (\text{Expected revenue from agent } i) \\
&\geq \sum_{i\in[n]} \left(\frac{1}{2}\right) \cdot Z_i(q_i) \\
&= \frac{1}{2} \sum_{i\in[n]} Z_i(q_i) \\
&\geq \frac{1}{4}(\text{EA-OPT}) \\
&\geq \frac{1}{4}(\text{OPT})
\end{aligned}
$$

■

4.3 Prophet Inequality

Consider the following online game. There are n boxes each of which has a reward inside it. The reward inside box i is the random variable X_i which is drawn from a known distribution F_i. We are allowed to open the boxes one at a time, observe the reward inside, and choose to either accept or reject the box. If we choose to accept the reward, the game ends immediately, and we receive the reward inside the box. If we reject the box, then we move on to another box, and the game continues.

In this game, a prophet is an entity that knows the instantiations of all the rewards. So, the prophet can always accept the box with the highest reward. Therefore, the expected reward of the prophet is $\mathbb{E}[\max X_i]$.

Theorem 4.3.1 (Prophet Inequality) *There exists an algorithm which can acheive $\frac{1}{2}$ of the prophet's expected reward by picking a simple threshold and accepting the first box whose reward crosses it.*

Proof: Let t be the quantity such that

$$\Pr[\exists\ i : X_i \geq t] = \frac{1}{2}$$

In other words, $t = \text{median}(\max_i X_i)$. Now, note that

$$
\begin{aligned}
\text{OPT} &= \mathbb{E}[\max_i X_i] \\
&= t + \mathbb{E}[\max_i X_i - t] \\
&\leq t + \mathbb{E}[\max_i (X_i)^+ - t]
\end{aligned}
$$

where $a^+ := \max\{a, 0\}$.

$$\begin{aligned} \therefore \text{OPT} &\le t + \mathbb{E}[\max_i (X_i - t)^+] \\ &\le t + \sum_{i \in [n]} \mathbb{E}[(X_i - t)^+] \end{aligned}$$

Here, the inequality $\mathbb{E}[\max_i (X_i)^+ - t] \le \mathbb{E}[\max_i (X_i - t)^+]$ follows from the fact that

$$\begin{aligned} \max_i (X_i, 0) - t &\le \max_i (X_i, t) - t \\ &= \max_i (X_i - t, 0) \end{aligned}$$

Now, let us consider the performance of the threshold algorithm. Note that by choice of t, the algorithm will receive a reward of t with probability $\frac{1}{2}$. Moreover, the reward may exceed t, if the box has a reward $X_i > t$. Therefore,

$$\text{ALG} = \frac{1}{2} t + \sum_{i \in [n]} \Pr[\text{ALG did not accept before box } i] \cdot \mathbb{E}[(X_i - t)^+]$$

The above equality follows from the fact that the random variable $\mathbb{1}_{\{\text{ALG did not accept before box } i\}}$ and the random variable $(X_i - t)^+$ are independent.

Notice that

$$\begin{aligned} \Pr[\text{ALG did not accept before box } i] &= \Pr[X_1 < t, X_2 < t, \ldots, X_{i-1} < t] \\ &\ge \Pr[X_1 < t, X_2 < t, \ldots, X_n < t] \\ &= \Pr[\max_i X_i < t] \\ &= \frac{1}{2} \end{aligned}$$

by our choice of t. Therefore,

$$\begin{aligned} \text{ALG} &\ge \frac{1}{2} t + \sum_{i \in [n]} \left(\frac{1}{2} \mathbb{E}[(X_i - t)^+] \right) \\ &= \frac{1}{2} \left(t + \sum_{i \in [n]} \mathbb{E}[(X_i - t)^+] \right) \\ &\ge \frac{1}{2} (\text{OPT}) \end{aligned}$$

■

This result suggests an alternate order-oblivious posted price mechanism in the single-item regular distribution setting. We can view the n boxes as the n agents, and the reward X_i as being the virtual value of each agent. Therefore, the above theorem suggests that if we sold the item to the first buyer whose virtual value exceeds the threshold t, then, we acheive a 2-approximation to the optimal mechanism. When the distributions are regular, the virtual value exceeds the threshold t precisely when the value exceeds $\phi_i^{-1}(t)$. Formally, the mechanism is as follows:

OPM Mechanism using Prophet inequality:

1. Calculate the threshold $t = \text{median}(\max_i X_i)$.
2. Set a posted price of $\phi_i^{-1}(t)$ for each agent i.

When the distributions are non-regular, then it is easily seen that it is sufficient to use the ironed virtual surplus $\overline{\phi_i}$ instead of the virtual surplus ϕ_i to obtain the same approximation ratio.

Corollary 4.3.2 *In a single item auction with n agents each having non-iid valuations v_i drawn from distributions F_i, the above mechanism gets a 2-approximation to the optimal mechanism.*

In this mechanism t will be strictly positive in most cases. This increased posted price of $\phi_i^{-1}(t)$ over the monopoly reserve price of $\phi_i^{-1}(0)$ captures the competition between the agents.

The prophet inequality mechanism above uses different threhsold prices and gets a better approximation than the ex-ante relaxation mechanism that we saw before. However, as we shall see later in the course, there are several settings where the prophet inequality mechanism will not be applicable, where the ex-ante relaxation will still be useful.

Additionally, note that both mechanisms we have discussed so far have non-anonymous posted prices, since the price of the item is different for each agent.

4.4 Approximation through increased competition

4.4.1 Bulow-Klemperer theorem

Let us recall the setting with n iid agents with regular values. We know that the optimal mechanism in this setting is the Vickrey auction with a reserve price. Intuitively, we can see that when n is really large, many agents will have a value larger than the reserve price, in which case, the auction behaves like a regular Vickrey auction. This suggests that when n is large, the reserve price is not very important to the outcome of the auction. The Bulow-Klemperer theorem formalizes this intuition.

Theorem 4.4.1 (Bulow-Klemperer theorem) *In a single item auction with the agents having iid valuations drawn from a regular distribution F, the revenue of the Vickrey auction over $n+1$ agents (with no reserve price) is at least as large as the revenue of the optimal auction over n agents.*

Proof: The Vickrey auction over $n+1$ buyers is the mechanism that sells the item to the agent with the highest virtual value (even when this virtual value is negative). Therefore, the Vickrey auction is the optimal mechanism among all the mechanisms that always sell the item.

We now construct an alternate mechanism over the $n+1$ buyers. We run the optimal mechanism over the first n buyers, and if the item goes unsold, we simply give it to the $(n+1)^{\text{th}}$ agent for free. Notice that this is a mechanism over the $n+1$ items which always sells the item. Moreover, the revenue of this mechanism is preceisely $\text{OPT}(n)$. Therefore,

$$\text{OPT}(n) \leq \text{Vickrey}(n+1)$$

■

This shows that when the agents have their values drawn from a iid regular distributions, the auctioneer is better off recruiting more agents to the auction, than trying to compute the reserve price for the auction.

4.4.2 Single sample mechanism

Consider the setting where we have only one buyer whose value v is drawn from a regular distribution F. In this setting, we already know that the optimal mechanism is the posted price mechanism with the monopoly reserve price $\phi^{-1}(0)$. Consider the following mechanism.

Single sample mechanism:

1. Draw $p \sim F$.
2. Offer the price p to the agent.

Theorem 4.4.2 *In a single item auction with a single agent whose value v is drawn from a regular distribution F, the single sample mechanism gets a 2-approximation to the optimal mechanism.*

Proof: This can be proven geometrically. First, note that the expected revenue from the single sample mechanism is given as

$$\mathbb{E}_{q \sim \text{Unif}[0,1]}\left[R(q)\right] = \int_0^1 R(q)dq$$

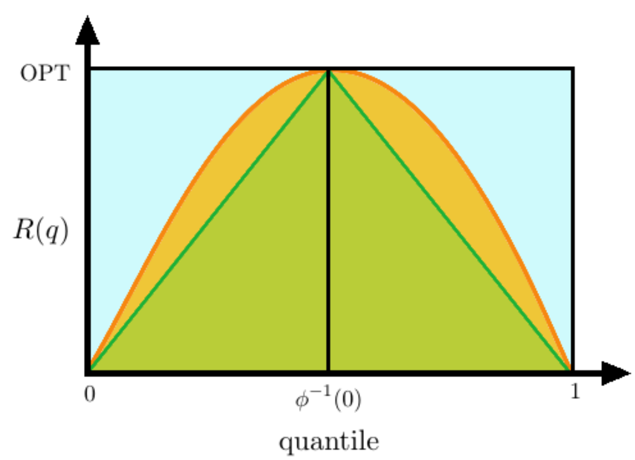

Figure 4.4.2: The revenue curve of a regular distribution

Therefore, the revenue from the single sample mechanism is the area under the orange revenue curve in the figure. Now, we know that that the revenue of the optimal mechanism is OPT and this is equal to the area under the blue rectangle. Since F is a regular distribution, we know that $R(q)$ is a concave function. Therefore, the area under the revenue curve is at least as large as the area under the green triangle. Now, note that the triangle has a base 1 and height OPT. So, the are of the triangle is $\frac{1}{2}$OPT.

$$\therefore \text{ALG} = \text{area under revenue curve} \geq \text{area under triangle} = \frac{1}{2}\text{OPT}$$

■

We can use this result to prove the following.

Theorem 4.4.3 *If there are n non-iid regular agents with values $v_1, \ldots, v_n$ drawn from distributions $F_1, \ldots, F_n$, then if we double the agents, so that we have n more agents with values $v'_1, \ldots, v'_n$ drawn from the same distributions $F_1, \ldots, F_n$, the Vickrey auction over the $2n$ agents gets at least $\frac{1}{2}$ of the revenue as the optimal mechanism over the n original agents.*

Intuitively, this can be explained by the fact that the value v'_i of the double acts as a single sample drawn from the distribution F_i. The formal proof of this theorem will be covered in the next lecture.

References

[1] Hartline, Jason D. "Mechanism design and approximation." Book draft. October 122 (2013).

[2] Roughgarden, Tim, and Qiqi Yan. "Revenue Maximization with a Single Sample." (2010).

[3] Hartline, Jason D., and Tim Roughgarden. "Simple versus optimal mechanisms." In Proceedings of the 10th ACM conference on Electronic commerce, pp. 225-234. 2009.