ORIE 6180: Design of Online Marketplaces **Spring 2016**

Lecture 6 — February 22

Lecturer: Sid Banerjee *Scribe: Venus Lo, David Eckman*

6.1 Overview of the last lecture

In the last lecture, we wanted to find a DSIC auction that maximizes expected revenue. We showed that if bidder i has valuation $v_i \sim F_i$ where distributions F_i are independent and known, then we can allocate and price using the VCG mechanisms on virtual valuations:

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

When $\phi_i(\cdot)$ is monotone increasing, we say that $F_i(\cdot)$ is *regular*. In order for this auction $\mathcal{A} = (x, p)$ to be DSIC, we require that $\phi_i(\cdot)$ be monotone increasing for all i. Because the auction is DSIC, the bids are equal to the bidders' valuations, i.e., $b_i = v_i$. We can then set prices according to Myerson's Lemma. Lastly,

$$\mathbb{E}[\text{Rev}(\mathcal{A})] = \mathbb{E}_{F_1,\ldots,F_n}\left[\sum_{i=1}^{n} \phi_i(v_i)x_i(v)\right]$$

Note that only bidders with virtual valuations $\phi_i(v_i) \geq 0$ are considered in the auction and that there is a reserve price of $\phi_i^{-1}(0)$ if i wins the auction.

6.2 Overview of this lecture

In this lecture, we will see that VCG mechanism with virtual valuation does not necessarily make intuitive sense, so it is hard to justify its use. Instead, we want to consider simple auctions where we can lower-bound the expected revenue by some fraction of the optimal solution. In other words, we want to look at approximately optimal mechanisms.

6.3 Review of expected revenue-maximizing auctions

Consider a single-item auction with n bidders and regular valuations. To maximize expected revenue, the optimal allocation rule is to pick the bidder i with the highest virtual valuation $\phi_i(v_i)$ and charge a price $p_i = \max\{\phi_i^{-1}(0), \max_{j \neq i} \phi_i^{-1}(\phi_j(v_j))\}$. If the valuation distributions F_i are iid, so that $v_i \sim F$ for a common F, then the monopoly reserve price is $r = \phi^{-1}(0)$. This setup is equivalent to running a 2^{nd}-price auction with reserve price r.

As an example, suppose $F_1 \sim \text{Unif}[0,1]$ and $F_2 \sim \text{Unif}[0,2]$. Straightforward calculations show that

$$\phi_1(v_1) = 2v_1 - 1 \quad \text{and} \quad \phi_2(v_2) = 2v_2 - 2.$$

The item is then allocated to the bidder with the highest virtual valuation, $\max\{\phi_1(v_1), \phi_2(v_2)\}$.

This type of auction may lead to bizarre results; for example if $x_1 = 1$, then

$$p_1 = \max\{\phi_1^{-1}(0), \phi_1^{-1}(\phi_2(v_2))\} = \max\{1/2, v_2 - 1/2\}.$$

If $v_1 = 3/4$ and $v_2 = 1$, then the item would be given to bidder 1 rather than bidder 2, and bidder 1 would pay 1/2. Therefore in some cases the item may be awarded to bidder 1 even though bidder 2 has a higher real valuation.

For practical settings, we want simple auctions. We will attempt to relax our objective of optimizing revenue to approximately optimizing revenue; see [1] Chapter 4-5 for more details.

High level approaches:

- Use knowledge of F_i to design simpler auctions
- Resource augmentation (ensure more bidders)
- Learn F_i from samples (kind of like bandit problems)

A typical simple auction that we will consider is a posted-price auction, one in which we set discriminatory reserve prices.

6.4 Prophet inequality

Consider an n-stage game in which at stage i, a gambler is offered a prize $\Pi_i \sim F_i$ and must decide whether to take the prize or move on to the next stage. Assume that the F_i's are independent and known and that the Π_i's are only revealed at stage i. The objective of this optimal stopping problem is to maximize the gambler's expected reward.

This problem can be solved using dynamic programming. It can be shown that the optimal policy is to set thresholds $\{t_i\}$, and stop the game when $\Pi_i \geq t_i$. However, we prefer to get a simpler, approximate solution with a single threshold t for all stages. That is, the single-threshold strategy $\mathcal{A}_t$ is to accept the first prize such that $\Pi_i \geq t$. Clearly this strategy is suboptimal because we may end up selecting no prize when we would do no worse by accepting the last prize in those cases.

The next result, commonly known as the "prophet inequality", shows that there exists a single-threshold policy whose expected reward is at least half of the expected reward of the oracle setting in which all rewards are known in advance.

Theorem 6.1. *(Samuel-Cohn) For every sequence $\{F_i\}_{i=1}^n$, there exists a threshold policy $\mathcal{A}_t$ that achieves*

$$\mathbb{E}[\mathrm{Rew}(\mathcal{A}_t)] \geq \frac{1}{2}\mathbb{E}[\max_i \Pi_i]. \tag{6.1}$$

Proof: Let $q(t) = \mathbb{P}[\text{No prize is picked by } \mathcal{A}_t]$. Then using the law of total probability, the left-hand side of (6.1) is bounded below by

$$\begin{aligned}
\mathbb{E}[\mathrm{Rew}(\mathcal{A}_t)] &= t(1-q(t)) + \mathbb{E}[\Pi_i - t : \Pi_i \text{ is first prize s.t. } \Pi_i \geq t] \\
&\geq t(1-q(t)) + \sum_{i=1}^n \mathbb{E}[\Pi_i - t \,|\, \Pi_i \geq t,\, \Pi_j < t\ \forall j \neq i]\, \mathbb{P}[\Pi_i \geq t]\, \mathbb{P}[\Pi_j < t\ \forall j \neq i] \\
&\geq t(1-q(t)) + \sum_{i=1}^n \mathbb{E}[(\Pi_i - t)^+] q(t).
\end{aligned}$$

The right-hand side of (6.1) is bounded above by

$$\begin{aligned}
\mathbb{E}[\max_i \Pi_i] &= \mathbb{E}[t + \max_i(\Pi_i - t)] \\
&\leq t + \mathbb{E}[\max_i(\Pi_i - t)^+] \\
&\leq t + \sum_{i=1}^n \mathbb{E}[(\Pi_i - t)^+].
\end{aligned}$$

We can then choose t such that $q(t) = 1/2$. Substituting this choice of t gives the desired inequality. □

Is it always true that there exists t where $q(t) = 1/2$? If q is continuous, then the answer is clearly yes. For the case where q is discrete, see [1] Chapter 4 for details on how to use t.

6.4.1 Discriminatory pricing auctions

Returning to single-item auctions, suppose that $v_i \sim F_i$ are independent where the distributions F_i are regular. We would like to determine posted prices r_i for each bidder.

By choosing a value $\hat{r}$ such that $\mathbb{P}[\max_i \phi_i(v_i)^+ \geq \hat{r}] = 1/2$ and setting a reserve price of $r_i = \phi_i^{-1}(\hat{r})$, we can appeal to Theorem 6.1 where we view the customers as the n stages where the ordering no longer matters. To allocate the item, we do a lottery among the bidders i satisfying $\phi_i(v_i) \geq \hat{r}$ and charge the winning bidder its reserve price of r_i.

In fact, we don't need to do a lottery to allocate the item. We can use any allocation rule that does not look at the bids amongst bidders with $\phi_i(v_i) \geq \hat{r}$, such as a priority ordering on the bidders.

6.4.2 Prior independent auctions

Assume $v_i \sim F$ iid. Then even if we don't know F, Theorem 6.2 states that we can still upper-bound the expected revenue of an optimal n-bidder auction ($\mathbb{E}[\text{Rev}(OPT_n)]$) by the expected revenue of a $n+1$-bidder VCG auction ($\mathbb{E}[\text{Rev}(VCG_{n+1})]$).

Theorem 6.2. *(Bulow-Klemperer) For a single-item auction with n agents and iid values $v_i \sim F$, F regular,*

$$\mathbb{E}[\text{Rev}(OPT_n)] \leq \mathbb{E}[\text{Rev}(VCG_{n+1})]. \tag{6.2}$$

Proof: Define a third mechanism $\mathcal{A}_{n+1}$ on $n+1$ bidders:

1. Run OPT_n on n bidders.
2. If item is not awarded, then give it to the $(n+1)^{\text{st}}$ bidder for free.

Clearly $\mathbb{E}[\text{Rev}(\mathcal{A}_{n+1})] = \mathbb{E}[\text{Rev}(OPT_n)]$.

Claim: Among all DSIC auctions on $n+1$ bidders which always allocate the item, VCG_{n+1} has the highest expected revenue. Recall that VCG with virtual valuations maximizes expected revenue. Since all bidders' valuations are iid, we have a single $\phi(\cdot)$. Furthermore, because the item must be allocated, there is no reserve price $\phi^{-1}(0)$. By monotonicity of $\phi(\cdot)$, this is exactly the same as running VCG on actual valuations.

The claim implies that $\mathbb{E}[\text{Rev}(VCG_{n+1})] \geq \mathbb{E}[\text{Rev}(\mathcal{A}_{n+1})]$, which completes the proof. □

The above result shows that the value (in terms of revenue increase) to an auctioneer of optimizing over auction formats is essentially the same as that of adding another bidder. This can be easily converted to the following direct comparison between welfare and revenue maximizing auctions for n bidders:

Corollary 6.3. *For a single-item auction with n agents and iid values $v_i \sim F$, F regular:*

$$\mathbb{E}[\text{Rev}(VCG_n)] \geq \left(\frac{n-1}{n}\right) \mathbb{E}[\text{Rev}(OPT_n)]$$

Proof: Let $Y_{2,n}$ denote the second-largest of n random variables drawn i.i.d from any distribution F (with associated pdf f) taking non-negative values. By definition, $\mathbb{E}[\text{Rev}(VCG_{n+1})] = \mathbb{E}[Y_{2,n}]$. Also, observe that $\mathbb{P}[Y_{2,n} \leq x] = F^n(x) + nF^{n-1}(x)(1-F(x))$, and thus:

$$\begin{aligned}
\mathbb{E}[Y_n] &= \sum_{x=0}^{\infty} x \cdot n(n-1)F^{n-2}(x)(1-F(x))f(x)dx \\
&\leq \frac{n}{n-1} \sum_{x=0}^{\infty} x \cdot (n-1)F^{n-2}(x)f(x)dx = \frac{n}{n-1}\mathbb{E}[Y_{n-1}]
\end{aligned}$$

Now using the Bulow-Klemperer theorem (Theorem 6.2), we get the result. □

Bibliography

[1] Hartline, J. D. Mechanism design and approximation.