COMP/MATH 533: Algorithmic Game Theory	September 22, 2014
Lecture 6	
Lecturer: Yang Cai	*Scribe: Jean-Benoit Chasles and Simon Labute*

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

In this lecture we look into some weaknesses of Myerson's auction and show how to design an approximately optimal auction with simpler and more transparent allocation and payment rules.

1 Examples of Myerson's Auction

1.1 Reminder of revenue of Myerson's Auction

We present a reminder of the result that in any single dimensional environment, we have that the expected revenue is

$$\mathbb{E}_{v\sim F}[\sum_i p_i(v)] = \mathbb{E}_{v\sim F}[\sum_i x_i(v)\varphi_i(v_i)]$$

Where $\varphi_i(v_i) := v_i - (1 - F_i(v_i))/f_i(v_i)$ which is called bidder i's virtual value and f_i is the density function for F_i. To optimize revenue, we use the virtual welfare maximizing allocation rule

$$x(v) := \text{argmax}_{x\in X} \sum_i x_i(v)\varphi_i(v_i)$$

1.2 Single item auction with values drawn from regular i.i.d distribution

Since all virtual value functions $\varphi^{-1}()$ are identical and monotone, the highest bidder has the highest virtual value and the optimal auction is the Vickrey auction (seen in previous lectures). This result is a simple mechanism for both bidders and the aucitoneer. However, for the non i.i.d case, the optimal auction which arises from Myerson's result may be too complex for practice as seen next.

Example 1. *2 bidders, v_1 uniform in [0,1]. v_2 uniform in [0,100]. $\varphi_1(v_1) = 2v_1 - 1, \varphi_2(v_2) = 2v_2 - 100$*
Under these settings, the auction which optimizes social welfare is as follows:
- Case 1: $v_1 > 1/2, v_2 < 50$, we sell to bidder 1 at price 1/2.
- Case 2: $v_1 < 1/2, v_2 > 50$, we sell to bidder 2 at price 50.
- Case 3: $0 < 2v_1 - 1 < 2v_2 - 100$, we sell to bidder 2 at price $(99 + 2v_1)/2$.
- Case 4: $0 < 2v_2 - 100 < 2v_1 - 1$, we sell to bidder 1 at price $(2v_2 - 99)/2$.

In Case 1, the allocation rule is quite counter-intuitive, as the item is sold to bidder 1, albeit his bid might actually be much lower than bidder 2's. Also the payments in Case 3 and 4 are not easy to explain to the bidders, as such numbers are not the reserve price or anyone else's bid. This shows that the optimal auction could be very complex, and the complexity of the mechanism causes the loss of interest for participating due to the high barrier to playing. Aside from the complexity, we also have the added issue that to design the optimal Myerson auction, we must be certain in advance of the underlying bidder's distributions, a situation which is rare in practice. This provides motivation for simpler mechanism design with approximately optimal revenue.

To design a simple nearly-revenue-optimal auction, we borrow tools from the optimal stopping theory — the Prophet Inequality.

2 Prophet Inequality

Consider the following game. For n stages (you are aware of how many stages in advance), you are offered a non-negative reward π_i where each reward is drawn from some known distributions G_i independently. Each reward π_i is revealed only at stage i. You decide to either accept the offered reward at stage i, or to discard that one and continue on. It is unclear how one should play this game. Imagine there is a prophet who sees all the realization of the π_i's before the game starts. He will certainly pick the largest π_i. The following theorem states that even though you can't see all the realizations of the rewards, there is a simple strategy which guarantees a half of the reward of a prophet!

Theorem 1 (Prophet inequality)**.** *There exists a strategy such that the expected reward* $\geq 1/2 \cdot \mathbb{E}[\max_i \pi_i]$ *Specifically, a threshold strategy, in which the first reward that is above the threshold t will be accepted, is sufficient to achieve this result.*

Proof: We leave t to be set at a later point.

Definition 1. *Given a random variable Z, we define Z^+ as* $\max(0, Z)$*.*

we also let $q(t) = \Pr[\pi_i < t,\ \forall i]$

We proceed in two steps to showing the result. First, we prove a lower bound on the expected reward of the threshold strategy and then we prove an upper bound on the reward of the prophet.

Lemma 1. *The expected reward of the threshold strategy is at least*

$$t(1 - q(t)) + q(t) \sum_i E[(\pi_i - t)^+]$$

.

Proof: : How do we give a lower bound on the expected reward? An obvious lower bound is $t \cdot (1 - q(t))$ as with probability $1 - q(t)$ some reward will be higher than the threshold. However, this could be way lower than what the reward really is. Imagine that we set $t = 100$, when some reward is above 100, it could be 120, 500 or even $1,000,000$, and you actually get the reward not just t. If there is only a single reward $\pi_i \geq t$, it is clear that we will receive π_i, and it is easy to write the expected reward for this event $t + \mathbb{E}[\pi_i - t | \pi_i \geq t, \forall_{j \neq i} \pi_j < t]$. When there are multiple rewards above t, we get the first one, and it is not easy to write a simple formula to represent this. So we relax our bound and say if there are more than one reward above t, we count our reward as t, which is an obvious lower bound.

So we have obtained an lower bound for our expected reward

$$\begin{aligned}
&t(1 - q(t)) + \sum_i \Pr[\pi_i \geq t, \forall_{j \neq i} \pi_j < t] \cdot \mathbb{E}[\pi_i - t | \pi_i \geq t, \forall_{j \neq i} \pi_j < t] \\
=&t(1 - q(t)) + \sum_i \Pr[\pi_i \geq t] \cdot \Pr[\forall_{j \neq i} \pi_j < t] \cdot \mathbb{E}[\pi_i - t | \pi_i \geq t] \quad (independence\ of\ the\ \pi_i's) \\
=&t(1 - q(t)) + \sum_i \mathbb{E}[(\pi_i - t)^+] \cdot \Pr[\forall_{j \neq i} \pi_j < t] \quad (the\ definition\ of\ (\pi_i - t)^+) \\
\geq&t(1 - q(t)) + q(t) * \sum_i \mathbb{E}[(\pi_i - t)^+] \quad (\Pr[\forall_{j \neq i} \pi_j < t] \leq q(t))
\end{aligned}$$

□

Next, we show a simple upper bound for the expected reward of the prophet.

Lemma 2. *we have that the expected reward is at most*

$$\mathbb{E}[\max_i \pi_i] \leq t + \sum_i \mathbb{E}[(\pi_i - t)^+]$$

Proof:

$$\begin{aligned}&\mathbb{E}[\max_i \pi_i]\\ =&\mathbb{E}[t+\max(\pi_i - t)] = t + \mathbb{E}[\max(\pi_i - t)]\\ \le& t + \mathbb{E}[\max(\pi_i - t)^+] \le t + \mathbb{E}[\sum_i (\pi_i - t)^+]\\ =& t + \sum_i \mathbb{E}[(\pi_i - t)^+]\end{aligned}$$

□

If we choose t such that $q(t) = 1/2$, then combining these two results, we see that the expected reward of the threshold is at least $1/2 \cdot \mathbb{E}[\max_i \pi_i]$. □

Remark 1. *We have actually proved a stronger version of this theorem. In our lower bound, we only count the reward as t when there are multiple rewards above t. If we change the threshold strategy to be — pick the smallest reward that is above the threshold, the same lower bound holds. Thus, the same proof can show the expected reward of this new strategy is still at least $1/2 \cdot \mathbb{E}[\max_i \pi_i]$. This stronger version of prophet inequality is what we will use in the design for simple nearly-optimal auctions.*

3 Simple Nearly-Optimal Auctions

Using this result, we revisit the case of a single item auction with bidders' value distributions being non i.i.d. We think of an auction as the game introduced above where the virtual value functions $\varphi_i(v_i)^+$ as the ith prize where G_i is the induced non-negative virtual value distribution from F_i. The prophet's expected reward in this game $\mathbb{E}_{v\sim F}[\max_i \varphi_i(v_i)^+]$ is exactly the optimal revenue in the corresponding single-item auction. If we can find an auction whose expected virtual welfare is at least a constant fraction of $\mathbb{E}_{v\sim F}[\max_i \varphi_i(v_i)^+]$, then we have an auction that achieves a constant fraction of the optimal revenue.

Consider the allocation rule:
1. choose t such that $\Pr[\max_i \varphi_i(v_i)^+ \ge t] = 1/2$
2. give the item to some bidder i with $\varphi_i(v_i) \ge t$ if there are ties, break them arbitrarily (subject to monotonicity).
By Prophet Inequality, any allocation rule which satisfies the above has

$$\mathbb{E}_{v\sim F}[\max \sum_i x_i(v)\varphi_i(v_i)] \ge 1/2 \cdot \mathbb{E}_{v\sim F}[\max_i \varphi_i(v_i)^+].$$

Therefore, any of these allocation rule can be turned to an auction that achieves at least half of the optimal revenue.

More specifically, we can select for instance the following allocation rule:
1. For each bidder i, set the reserve price $r_i = \varphi_i^{-1}(t)$ using the t defined above.
2. Give the item to the highest bidder that meets his or her reserve price, if any.

As for the payment rule, it is simply the maximum between the winner's reserve price and the second highest bid (that meets the reserve). We see here a much simpler mechanism which achieves at least half optimal revenue.

This mechanism is much simpler than Myerson's auction when the distributions are i.i.d.. An even simpler auction would use the same reserve price for every bidder. An important open problem asks the following:
Open Problem: If we are restricted to use the same reserve price, what is the best approximation ratio we can achieve for revenue. We know the answer is between 1/2 and 1/4.