CS 580: Algorithmic Game Theory
Lecture 11: Simple vs Optimal Auctions

Vasilis Livanos*

1 Impracticality of Optimal Auctions

We saw in the last lecture that for every DSIC auction (x, p) for a single-parameter environment with valuation distributions $D_1, \ldots, D_n$, the expected revenue equals the expected virtual welfare. This is one of the most fundamental results in auction theory, as it allows us to study only auctions where the objective is to maximize the social welfare. We also obtained a corollary of this result; the optimal single-item auction with bidders whose valuations are drawn i.i.d. from a regular distribution D is the Vickrey auction, augmented with the reserve price $\phi^{-1}(0)$, where ϕ is the virtual welfare function, defined as

$$\phi(v) = v - \frac{1 - F_D(v)}{f_D(v)},$$

where F_D and f_D is the cdf and pdf of D respectively. We saw that, when D is regular, ϕ is strictly increasing, the virtual welfare-maximizing allocation rule is monotone and, for suitable payments, it maximizes the expected revenue over all DSIC auctions.

What if one wanted to extend this result? One way would be to relax the assumption that the bidders' valuations come from the same distribution D; another would be to consider selling more than one item. We will take a look at both of these later. Yet another, which is the focus of the next chapter is to assume we do not have access to the distribution D.

2 The Bulow-Klemperer Theorem

How does one design an auction when they do not have any information on the distribution of the bidders' valuations? The answer given by Bulow and Klemperer in their celebrated theorem can be quite surprising. In this first setting we will examine, we assume that, while D exists and the bidders' valuations are drawn i.i.d. from it, we do not know D. Therefore, we cannot use any distributional information in the *design* of our auction, but we can use it in its *analysis*. This is a well-motivated assumption as, in some applications either the bidders' valuations might be changing very rapidly or we might simply not have enough data to infer the underlying distribution.

As it turns out, this assumption does not make the problem overwhelmingly difficult; in fact, to compete and actually outperform the optimal auction with n bidders, it suffices to incentivize one extra bidder to join in, as the following theorem by Bulow and Klemperer [1] shows.

Theorem 1. *Let D be a regular distribution and n a positive integer. Then:*

$$\mathop{\mathbb{E}}_{v_1,\ldots,v_{n+1}\sim D}[Rev(VA_{n+1})] \geq \mathop{\mathbb{E}}_{v_1,\ldots,v_n\sim F}[Rev(OPT_{D,n})], \tag{1}$$

where VA_{n+1} and $OPT_{D,n}$ denote the Vickrey auction on $n+1$ bidders and the optimal auction for D on n bidders, respectively.

Notice that the left-hand side auction is the Vickrey auction with no reserve price.

Proof (of Theorem 1). We define an "intermediary" auction $\mathcal{A}$ that will allow us to compare the two sides of (1). $\mathcal{A}$ considers $n+1$ bidders in the following way: it first simulates the optimal auction on the first n

* University of Illinois Urbana-Champaign, Urbana IL 61801, USA livanos3@illinois.edu

bidders. If the optimal auction does not allocate the item to any of the first n bidders, then $\mathcal{A}$ allocates the item to bidder $n+1$ for free.

First, notice that $\mathcal{A}$'s expected revenue is exactly that of the optimal auction with n bidders. Also, $\mathcal{A}$ always allocates the item to one of the bidders. Thus, if we can prove that the Vickrey auction maximizes the expected revenue over all auctions that are guaranteed to allocate the item, then we are done.

Consider the optimal auction that always allocates the item and denote it by $\mathcal{B}$. Notice that, by the equivalence of expected revenue and expected virtual welfare, $\mathcal{B}$ always awards the item to the bidder with the highest virtual valuation, even if this virtual valuation is negative. On the other hand, the Vickrey auction awards the item to the bidder with the highest valuation. However, since the bidders' valuations are i.i.d. draws from D, all bidders share the same virtual valuation function ϕ. Furthermore, since D is a regular distribution, ϕ is increasing. Thus, the bidder with the highest virtual valuation always has the highest valuation too.

Therefore, the Vickrey auction (with $n+1$ bidders) has expected revenue at least that of every auction that always allocates the item, including $\mathcal{A}$, and therefore its expected revenue is at least that of OPT_D (with n bidders).

As a corollary, this implies that the expected revenue of the Vickrey auction with n bidders and no reserve price is a $(1 - 1/n)$-approximation to that of the optimal auction with n bidders. The theorem also implies the following: it's more important for an auction to be competitive than for its design to be optimal. In other words, one should strive to have more participants involved, even if it means they run a simpler auction at the end of the day.

3 The Simplest Auction: Grocer's Auction

In the previous section, we relaxed the optimality of our auction by assuming we did not have access to the common distribution D from which the bidders' valuations are drawn. In this section, we will take a different approach. Instead of assuming there exists a single, unknown distribution D, we will assume that there exist n *known* distributions $D_1, \ldots, D_n$ and the valuation v_i of each bidder i is drawn from D_i, independently from the other bidders. This, unfortunately, makes the optimal auction much more complicated. In fact, it is even possible in cases that the highest bidder does not even win the optimal auction!

Obviously, such an auction would not be implementable in practice; for the system to work, the participants must be able to understand it and participate in good faith. For this reason, we will strive to design simpler auctions which are more practical than the optimal auction. Therefore, we necessarily will have to relax the guarantee of optimality and instead settle for auctions which are approximately optimal.

What *simple* auction would work reasonably well in such a setting? Thankfully, one does not need to look very far! Drawing inspiration from real life, we describe a very simple auction, which we call the *grocer's auction*.

In the grocer's auction, the grocer, who is the mechanism designer, decides on a single price to set for the item, which will be offered to all buyers. Notice that we refer to them now as buyers instead of bidders since the grocer never collects any bids. Also notice that the same price is offered to everyone, i.e. the price is *anonymous*. After that, the buyers arrive at the store one by one, in some order, and each buyer sees the item and its price, if it is still available. Then, a buyer decides to buy the item if and only if their own true valuation exceeds the price.

The question of course becomes: How do we select a good price p? Too high and no buyer will buy, too low and we will obtain very little revenue (or welfare, since the item will most likely be sold quickly to a buyer of low valuation). Also notice that we do not have any control over the order in which the buyers arrive; thus our mechanism needs to work for all arrival orders of the buyers or, in other words, for an *adversarial* ordering.

The above constraints led Hajiaghayi et al [3] and later Chawla et al [2] to reduce the above design problem to a neat problem from optimal stopping theory; the *prophet inequality*.

4 The Prophet Inequality

The classical *prophet inequality* setting is the following: We are given n independent random variables $X_1, X_2, \dots, X_n$ along with their distributions $D_1, D_2, \dots, D_n$, where $D_i : \mathbb{R}_{\geq 0} \to [0,1]$ for all i. We observe a realization z_i of each X_i sequentially, in an order determined by an adversary (the prophet). At step i, we have to immediately and irrevocably decide whether to select z_i, in which case the game ends, and we receive value z_i, or skip it, in which case we are not able to select it again.

Our objective is to select the highest possible value, and we compare our selection against the selection of a prophet, who sees all realizations at once and thus can always pick the maximum of all realizations. In other words, we compare against $\mathbb{E}[\max_i X_i]$.

Below, we denote $\{X_1, X_2, \dots, X_n\}$ and $\{D_1, D_2, \dots, D_n\}$ by $\boldsymbol{X}$ and $\boldsymbol{D}$, respectively.

4.1 Upper Bound

In 1977, Krengel and Sucheston [5,6] showed that there exists an online strategy which, on expectation over the random draws of the distributions, achieves half of the expected offline optimum value, i.e. $\frac{1}{2}\mathbb{E}[\max_i X_i]$. Later, Samuel-Cahn [8] showed that this ratio can be achieved by a simple threshold rule, in which one can set a threshold τ a priori, and whenever a realization exceeds that threshold, select it and stop.

This threshold value was set to the median of the distribution of $\max_i X_i$, i.e. τ is selected such that $\Pr[\max_i X_i > \tau] = \frac{1}{2}$. Later, Kleinberg and Weinberg [4] showed that setting $\tau = \frac{1}{2}\mathbb{E}[\max_i X_i]$ suffices as well. We provide both proofs below.

Theorem 2 **([8]).** *Consider an instance* $(\boldsymbol{X}, \boldsymbol{D})$ *of the prophet inequality setting, and let* τ *be the median of the distribution of* $\max_i X_i$*, i.e.* $\Pr[\max_i X_i > \tau] = \frac{1}{2}$*. An algorithm that selects the first value that exceeds* τ *receives value* V *such that*

$$\mathbb{E}[V] \geq \frac{1}{2}\mathbb{E}\left[\max_i X_i\right].$$

Proof. Let $X^* = \max_i X_i$, and $\mathcal{E}_i$ denote the event that our algorithms "reaches" i, i.e. has not accepted any of the first $i-1$ values. Since our algorithm accepts the first value above τ, we have

$$\begin{aligned}
\mathbb{E}[V] &= \tau \Pr[X^* \geq \tau] + \sum_{i=1}^{n} \Pr[\mathcal{E}_i]\,\mathbb{E}\left[(X_i - \tau)^+\right] \\
&= \frac{1}{2}\tau + \sum_{i=1}^{n} \Pr\left[\max_{1\leq j\leq i-1} X_j < \tau\right]\mathbb{E}\left[(X_i - \tau)^+\right] \\
&\geq \frac{1}{2}\tau + \sum_{i=1}^{n} \Pr[X^* < \tau]\,\mathbb{E}\left[(X_i - \tau)^+\right] \\
&\geq \frac{1}{2}\tau + \frac{1}{2}\sum_{i=1}^{n} \mathbb{E}\left[(X_i - \tau)^+\right] && \text{since } \Pr[X^* < \tau] = \frac{1}{2} \\
&\geq \frac{1}{2}\tau + \frac{1}{2}\mathbb{E}\left[(X^* - \tau)^+\right] \\
&\geq \frac{1}{2}\mathbb{E}[X^*].
\end{aligned}$$

The above proof is the simplest and the one that you will most likely find in other lecture notes or slides. However, in practice, it might be difficult to calculate the median of the distribution of $\max_i X_i$, whereas the expectation might be much simpler to compute (for example by sampling from the distributions and averaging). Therefore, the following threshold is a simpler one, even though the analysis is not as clean as the previous theorem.

Theorem 3 **([4]).** *Consider an instance* $(\boldsymbol{X}, \boldsymbol{D})$ *of the prophet inequality setting, and let* $\tau = \frac{1}{2}\mathbb{E}[\max_i X_i]$*. An algorithm that selects the first value that exceeds* τ *receives value* V *such that*

$$\mathbb{E}[V] \geq \frac{1}{2}\mathbb{E}\left[\max_i X_i\right].$$

Proof. Let $X^* = \max_i X_i$, and $\mathcal{E}_i$ denote the event that our algorithms "reaches" i, i.e. has not accepted any of the first $i-1$ values. Furthermore, let $p = \Pr[x^* \geq \tau]$. By the definition of expectation, we have

$$\mathbb{E}[V] = \int_0^\infty \Pr[V > u]\, du = \int_0^\tau \Pr[V > u]\, du + \int_\tau^\infty \Pr[V > u]\, du.$$

For the first term, notice that for all $u < \tau$

$$\Pr[V > u] = \Pr[V \geq \tau] = \Pr[X^* \geq \tau] = p.$$

since we only accept some value if and only if it is at least τ, and the probability we accept a value is exactly p. Therefore

$$\int_0^\tau \Pr[V > u]\, du \geq p\tau. \tag{2}$$

For the second term, notice that with probability $1-p$, the algorithm selects no random variable. Therefore $\Pr[\mathcal{E}_i] \geq 1-p$, which implies that, for any $u \geq \tau$

$$\Pr[V > u] = \sum_{i=1}^n \Pr[\mathcal{E}_i] \Pr[X_i > u] \geq (1-p) \sum_{i=1}^n \Pr[X_i > u] \geq (1-p) \Pr[X^* > u],$$

where the last inequality follows from a simple union bound. However, notice that

$$2\tau = \mathbb{E}[X^*] = \int_0^\tau \Pr[X^* > u]\, du + \int_\tau^\infty \Pr[X^* > u]\, du,$$

and also $\Pr[X^* > u] \leq 1$ for all u, which implies $\int_0^\tau \Pr[X^* > u]\, du \leq \tau$, and we obtain

$$\int_\tau^\infty \Pr[X^* > u]\, du \geq \tau.$$

Thus,

$$\int_0^\tau \Pr[V > u]\, du \geq (1-p) \int_0^\tau \Pr[X^* > u]\, du \geq (1-p)\tau. \tag{3}$$

Combining (2) and (3), we get

$$\mathbb{E}[V] \geq p\tau + (1-p)\tau = \tau = \frac{1}{2}\, \mathbb{E}[X^*].$$

Remark 1. Theorems 2 and 3 hold even against an "almighty" prophet who knows all the realizations and any random coins of the algorithm a priori.

Lower Bound As it turns out, this 1/2-competitive ratio is the best one can hope for. In other words, there exists an instance $(\boldsymbol{X}, \boldsymbol{D})$ in which no algorithm can obtain, in expectation, value better than $\frac{1}{2}\, \mathbb{E}\left[\max_i X_i\right]$.

Theorem 4. *There exists an instance $(\boldsymbol{X}, \boldsymbol{D})$ of the prophet inequality setting such that, for any online algorithm $\mathcal{A}$ which receives value V and any $\varepsilon > 0$, we have*

$$\mathbb{E}[V] < \left(\frac{1}{2} + \varepsilon\right) \mathbb{E}\left[\max_i X_i\right].$$

Proof. Consider the following instance $(\boldsymbol{X}, \boldsymbol{D})$ of the prophet inequality setting: We are given two random variables X_1 and X_2, where $X_1 = 1$ with probability 1 and $X_2 = 0$ with probability $1-\varepsilon'$ and $X_2 = \frac{1}{\varepsilon'}$ with probability ε', for some $\varepsilon' > 0$. Then, on expectation, any algorithm $\mathcal{A}$ receives the same value by either selecting X_1 or X_2, and that value is $\mathbb{E}[V] = 1$. However, the prophet, who can see the realization of X_2 offline, can select X_2 whenever $X_2 = \frac{1}{\varepsilon'}$. Therefore,

$$\mathbb{E}\left[\max_i X_i\right] = \varepsilon' \cdot \frac{1}{\varepsilon'} + (1-\varepsilon') \cdot 1 = 2 - \varepsilon',$$

and, for $\varepsilon' < \frac{2\varepsilon}{1/2+\varepsilon}$ we obtain

$$\mathbb{E}[V] = 1 < \left(\frac{1}{2} + \varepsilon\right)(2 - \varepsilon') = \left(\frac{1}{2} + \varepsilon\right) \mathbb{E}\left[\max_i X_i\right].$$

4.2 Connection with the Grocer's Auction

We saw two algorithms for the prophet inequality setting which achieve the tight 1/2-competitive ratio. Also, perhaps interestingly, both algorithms were (single) threshold algorithms; they decided on a threshold τ and selected the first realization above that threshold. In the reduction of the grocer's auction to the prophet inequality, this threshold corresponds to the price p that the grocer should set for the item. This will yield an approximation guarantee with respect to the optimal welfare and, via virtual valuations, to the optimal revenue.

Surprisingly, the optimal threshold of 1/2 can be achieved even if one has no full knowledge of $D_1, \ldots, D_n$ and only has access to a single sample from each of them [7]! In the next lecture, we will generalize prophet inequalities in a couple different directions and we will see that a single threshold does not suffice anymore for an optimal prophet inequality algorithm.

References

1. Bulow, J., Klemperer, P.: Auctions versus negotiations. American Economic Review **86**(1), 180–94 (1996), https://EconPapers.repec.org/RePEc:aea:aecrev:v:86:y:1996:i:1:p:180-94
2. Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: Proceedings of the Forty-Second ACM Symposium on Theory of Computing. p. 311–320. STOC '10, Association for Computing Machinery, New York, NY, USA (2010). https://doi.org/10.1145/1806689.1806733, https://doi.org/10.1145/1806689.1806733
3. Hajiaghayi, M.T., Kleinberg, R., Sandholm, T.: Automated online mechanism design and prophet inequalities. In: Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 1. pp. 58–65. AAAI'07, AAAI Press (2007), http://dl.acm.org/citation.cfm?id=1619645.1619656
4. Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proceedings of the forty-fourth annual ACM symposium on Theory of computing. pp. 123–136 (2012)
5. Krengel, U., Sucheston, L.: Semiamarts and finite values. Bull. Amer. Math. Soc. **83**(4), 745–747 (07 1977), https://projecteuclid.org:443/euclid.bams/1183538915
6. Krengel, U., Sucheston, L.: On semiamarts, amarts, and processes with finite value. Probability on Banach spaces **4**, 197–266 (1978)
7. Rubinstein, A., Wang, J.Z., Weinberg, S.M.: Optimal single-choice prophet inequalities from samples. In: Vidick, T. (ed.) 11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA. LIPIcs, vol. 151, pp. 60:1–60:10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). https://doi.org/10.4230/LIPIcs.ITCS.2020.60, https://doi.org/10.4230/LIPIcs.ITCS.2020.60
8. Samuel-Cahn, E.: Comparison of threshold stop rules and maximum for independent nonnegative random variables. the Annals of Probability pp. 1213–1216 (1984)