

CS364A: Algorithmic Game Theory

Lecture #6: Simple Near-Optimal Auctions*

Tim Roughgarden[†]

October 9, 2013

1 Optimal Auctions Can Be Complex

Last lecture we proved some of the most fundamental results in auction theory. To reiterate, for every DSIC auction (\mathbf{x}, \mathbf{p}) for a single-parameter environment with valuation distributions F_1, \dots, F_n , the expected revenue equals the expected virtual welfare:

$$\mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n p_i(\mathbf{v}) \right] = \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \varphi_i(v_i) \cdot x_i(\mathbf{v}) \right]. \quad (1)$$

Define the virtual welfare-maximizing allocation rule as the one that sets

$$\mathbf{x}(\mathbf{v}) := \operatorname{argmax}_X \sum_{i=1}^n \varphi_i(v_i) x_i(\mathbf{v})$$

for each input \mathbf{v} . If every F_i is regular, meaning that the corresponding virtual valuation function

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

is strictly increasing, then the virtual welfare-maximizing allocation rule is monotone and, after we define suitable payments, maximizes expected revenue over all DSIC auctions. This characterization of optimal auctions can be extended to irregular distributions, but this extension requires more work (see [3, Chapter 3]).

As a corollary of this general characterization, we noted that the optimal single-item auction with i.i.d. bidders and a regular distribution F is shockingly simple: it is simply the Vickrey auction, augmented with the reserve price $\varphi^{-1}(0)$. This is a true “killer application” of auction theory — it gives crisp, conceptually clean, and practically useful guidance to auction design.

*©2013, Tim Roughgarden.

[†]Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

The plot thickens if we make the problem a bit more complex. Consider again a single-item auction, but relax the assumption that bidders' valuation distributions are identical; they are still independent and regular. The optimal auction can get weird, and it does not generally resemble any auctions used in practice (see the exercises). For example, someone other than the highest bidder might win. The payment made by the winner seems impossible to explain to someone who hasn't studied virtual valuations — compare this to the i.i.d. case, where the optimal auction is simply eBay with a smartly chosen opening bid. This weirdness is inevitable if you are 100% confident in your model (i.e., the F_i 's) and you want every last cent of the maximum-possible expected revenue — there is no choice in which allocation or payment rule you can use.

Today we'll seek out auctions that are simpler, more practical, and more robust than the theoretically optimal auction. Since optimality requires complexity, we can only hope that our auctions are *approximately* optimal. This is the second time we've turned to approximation to escape a quandary posed by full optimality. In algorithm mechanism design (Lecture 4), we used approximation to recover computational tractability when the underlying welfare-maximization problem was NP-hard. Similarly, here we are rejecting the optimal auction because of its “complexity” and using approximation to recover relative simplicity. Unlike algorithmic mechanism design, where “low complexity” was precisely defined (as polynomial running time), here we leave terms like “simple”, “practical”, and “robust” largely undefined. Formulating definitions that capture these vague ideas and give useful guidance to the design of simple near-optimal auctions is an extremely important research question in the field.

The theory of Bayesian optimal auctions developed last lecture is part of the microeconomic canon, dating to 1981 [5]. By contrast, the research agendas outlined this lecture have been developed primarily over the past 5 years, mostly in the computer science literature. The classical theory from last lecture is the foundation on which this more recent work rests.

2 The Prophet Inequality

This section covers a fun result from optimal stopping theory. The next section uses it to design a relatively simple and provably near-optimal single-item auction for non-i.i.d. bidders.

Consider the following game, which has n stages. In stage i , you are offered a nonnegative prize π_i , drawn from a distribution G_i . You are told the distributions G_1, \dots, G_n in advance, and these distributions are independent. You are told the realization π_i only at stage i . After seeing π_i , you can either accept the prize and end the game, or discard the prize and proceed to the next stage. The decision's difficulty stems from the trade-off between the risk of accepting a reasonable prize early and then missing out later on a great one, and the risk of having to settle for a lousy prize in one of the final stages.

The Prophet Inequality, due to Samuel-Cahn [7], offers a simple strategy that does almost as well as a fully clairvoyant prophet.

Theorem 2.1 (Prophet Inequality) *For every sequence G_1, \dots, G_n of independent distributions, there is strategy that guarantees expected reward $\frac{1}{2}\mathbf{E}_\pi[\max_i \pi_i]$. In fact, there is such a threshold strategy t , which accepts prize i if and only if $\pi_i \geq t$.*

Proof: Let z^+ denote $\max\{z, 0\}$. Consider a threshold strategy with threshold t . It is difficult to compare directly the expected payoff of this strategy with the expected payoff of a prophet. So, the plan is to derive lower and upper bounds, respectively, on these two quantities that are easily to compare.

Let $q(t)$ denote the probability that the threshold strategy accepts no prize at all.¹ As t increases, the risk $q(t)$ increases but the average value of an accepted prize goes up.

What payoff does the t -threshold strategy obtain? With probability $q(t)$, zero, and with probability $1 - q(t)$, at least t . Let's improve our lower bound in the second case. If exactly one prize i satisfies $\pi_i \geq t$, then we should get "extra credit" of $\pi_i - t$ above and beyond our baseline payoff of t . If at least two prizes exceed the threshold, say i and j , then things are more complicated: our "extra credit" is either $v_i - t$ or $v_j - t$, depending on which of i, j belongs to the earlier stage. Let's be lazy and punt on this complication: when two or more prizes exceed the threshold, we'll only credit the baseline t to the strategy's payoff.

Formally, we have the following lower bound:

$$\begin{aligned} & \mathbf{E}[\text{payoff of } t\text{-threshold strategy}] \\ & \geq (1 - q(t)) \cdot t + \sum_{i=1}^n \mathbf{E}[\pi_i - t \mid \pi_i \geq t, \pi_j < t \forall j \neq i] \mathbf{Pr}[\pi_i \geq t] \cdot \mathbf{Pr}[\pi_j < t \forall j \neq i] \quad (2) \\ & = (1 - q(t)) \cdot t + \sum_{i=1}^n \underbrace{\mathbf{E}[\pi_i - t \mid \pi_i \geq t] \mathbf{Pr}[\pi_i \geq t]}_{=\mathbf{E}[(\pi_i - t)^+]} \cdot \underbrace{\mathbf{Pr}[\pi_j < t \forall j \neq i]}_{\geq q(t)} \quad (3) \end{aligned}$$

$$\geq (1 - q(t)) \cdot t + q(t) \sum_{i=1}^n \mathbf{E}[(\pi_i - t)^+], \quad (4)$$

where we use the independence of the G_i 's in (3) to factor the two probability terms and in (3) to drop the conditioning on the event that $\pi_j < t$ for every $j \neq i$. In (4), we use that $q(t) = \mathbf{Pr}[\pi_j < t \forall j] \leq \mathbf{Pr}[\pi_j < t \forall j \neq i]$.

Now we produce an upper bound on the prophet's expected payoff $\mathbf{E}[\max_i \pi_i]$ that is easy to compare to (4). The initial expression doesn't reference the strategy's threshold t , so we add and subtract it to derive

$$\begin{aligned} \mathbf{E}\left[\max_i \pi_i\right] &= \mathbf{E}\left[t + \max_i (\pi_i - t)\right] \\ &\leq t + \mathbf{E}\left[\max_i (\pi_i - t)^+\right] \\ &\leq t + \sum_{i=1}^n \mathbf{E}[(\pi_i - t)^+]. \quad (5) \end{aligned}$$

¹Note that discarding the final stage's prize is clearly suboptimal!

Comparing (4) and (5), we can set t so that $q(t) = \frac{1}{2}$ — i.e., there is a 50/50 chance of accepting a prize — to complete the proof.² ■

Our proof of Theorem 2.1 shows a stronger statement that is useful in the next section. Our lower bound (4) on the revenue of the t -threshold strategy only credits t units of value when at least two prizes exceed the threshold t ; only realizations in which exactly one prize exceeds the threshold contribute to the second, “extra credit” term in (4). This means that the guarantee of $\frac{1}{2}\mathbf{E}[\max_i \pi_i]$ applies even if, whenever there are multiple prizes above the threshold, the strategy somehow picks the worst (i.e., smallest) of these.

3 Simple Single-Item Auctions

We now return to our motivating example of a single-item auction with n bidders with valuations drawn from (not necessarily identical) regular distributions F_1, \dots, F_n . We use the Prophet Inequality to design a relatively simple, near-optimal auction.

The key idea is to regard the virtual valuation $\varphi_i(v_i)^+$ of a bidder, if nonnegative, as the i th prize. (G_i is then the corresponding distribution induced by F_i ; since the F_i ’s are independent, so are the G_i ’s.) To see an initial connection to the Prophet Inequality, note that the expected revenue of the optimal auction is $\mathbf{E}_{\mathbf{v}}[\sum_i \varphi_i(v_i)x_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\max_i \varphi_i(v_i)^+]$, precisely the expected value obtained by a prophet with prizes $\varphi_1(v_1)^+, \dots, \varphi_n(v_n)^+$.

Now can consider any allocation rule that has the following form:

- (1) Choose t such that $\mathbf{Pr}[\max_i \varphi_i(v_i)^+ \geq t] = \frac{1}{2}$.
- (2) Give the item to a bidder i with $\varphi_i(v_i) \geq t$, if any, breaking ties among multiple candidate winners arbitrarily (subject to monotonicity).

The strong form of the Prophet Inequality immediately implies that every auction with an allocation rule of the above type satisfies

$$\mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \varphi_i(v_i)^+ x_i(\mathbf{v}) \right] \geq \frac{1}{2} \mathbf{E}_{\mathbf{v}} \left[\max_i \varphi_i(v_i)^+ \right].$$

For example, here is a specific such allocation rule:

- (1) Set a reserve price $r_i = \varphi_i^{-1}(t)$ for each bidder i , with t defined as above.
- (2) Give the item to the highest bidder that meets its reserve (if any).

This auction first filters bidders using reserve prices, and then simply awards the item to the highest bidder remaining. This auction is qualitatively simpler than the optimal auction in two senses. First, the corresponding payment of the winning bidder is just the maximum of

²If there is no such t because of point masses in the G_i ’s, then a minor extension of the argument yields the same result (see Problem Set #2).

its reserve price and the highest bid by another bidder that meets its reserve price — thus, virtual valuation functions are only used to set reserve prices, and only the inverse virtual valuation of 0 matters. Second, the highest bidder wins, as long as it clears its reserve price.

This “simple” auction is more plausible to implement than an optimal auction, but an issue remains: the reserve prices are different for different bidders. Some real-world auctions use such non-anonymous reserve prices — in sponsored search auctions, “higher-quality” advertisers generally face lower reserve prices than lower-quality advertisers — but they are rare. On eBay, for example, you only get to set one opening bid, even if you know (from bidding histories, say) that the bidders are not i.i.d.

An interesting open research question is to understand how well the Vickrey auction with an anonymous reserve price (i.e., eBay) can approximate the optimal expected revenue in a single-item auction when bidders’ valuations are drawn from non-i.i.d. regular distributions. Partial results are known: there is such an auction that recovers at least 25% of the optimal revenue, and no such auction always recovers more than 50% of the optimal revenue [4].

More generally, designing simple auctions that provably approximate the optimal revenue has been a hot research topic for the past 5 years or so; see [3, Chapter 4] for a survey.

4 Prior-Independent Auctions and the Bulow-Klemperer Theorem

This section explores a different critique of the optimal auction approach developed last lecture: the valuation distributions F_1, \dots, F_n were assumed to be known to the seller in advance. In some applications, where there is lots of data and bidders’ preferences are not changing too rapidly, this is a reasonable assumption. But what if the seller does not know, or is not confident about, the valuation distributions? This is a relevant issue in “thin markets” where there is not much data, including keyword auctions for rarely used (but potentially valuable) search queries.

Removing advance knowledge of the distributions might seem to banish us to our single-bidder single-item quandary (Lecture 5) that motivated the Bayesian approach. The difference is that we will continue to assume that bidders’ valuations are drawn from distributions; it’s just that these distributions are unknown a priori. That is, we now use distributions in the *analysis* of auctions, but not in their *design*. The goal is to design an auction, whose description is independent of the underlying distributions, that performs almost as well as if the distributions were known in advance. This research agenda of designing good *prior-independent* auctions was articulated by Dhangwatnotai et al. [2] and has been an active topic over the past three years; see Hartline [3, Chapter 5] for a survey of the latest developments.

Today, we’ll cover a beautiful result from classical auction theory which is also an important precursor to the design of prior-independent auctions. The expected revenue of a Vickrey auction can obviously only be less than that of an optimal auction; yet the following result, due to Bulow and Klemperer [1], shows that this inequality reverses when the Vickrey auction’s environment is made slightly more competitive.

Theorem 4.1 (Bulow-Klemperer Theorem [1]) *Let F be a regular distribution and n a positive integer. Then:*

$$\mathbf{E}_{v_1, \dots, v_{n+1} \sim F}[\text{Rev}(VA) \text{ (} n+1 \text{ bidders)}] \geq \mathbf{E}_{v_1, \dots, v_n \sim F}[\text{Rev}(OPT_F) \text{ (} n \text{ bidders)}], \quad (6)$$

where VA and OPT_F denote the Vickrey auction and the optimal auction for F , respectively.³

Notice that the auction in the left-hand side of (6) — the Vickrey auction with no reserve — is “prior-independent,” meaning its description is independent of the underlying distribution F . The auction in the right-hand side of (6) depends on the underlying distribution F through its reserve price. In this sense, a single auction (the Vickrey auction) is simultaneously competitive with an infinite number of different optimal auctions, across all possible single-item environments with i.i.d. regular bidder valuations. The guarantee in Theorem 4.1 also implies that, in every such environment with $n \geq 2$ bidders, the expected revenue of the Vickrey auction is at least $\frac{n-1}{n}$ times that of an optimal auction (for the same number of bidders); see the Exercises.

The usual interpretation of the Bulow-Klemperer theorem, which also has anecdotal support in practice, is that extra competition is more important than getting the auction format just right. That is, invest your resources into getting more serious participants, rather than sharpening your knowledge of their preferences (of course, do both if you can!). See the Problems for more extensions and variations of the Bulow-Klemperer theorem.

Proof of Theorem 4.1: The two sides of (6) are tricky to compare directly, so for the analysis we define a fictitious auction \mathcal{A} to facilitate the comparison. This auction operates in the environment with $n+1$ bidders, as follows:

- (1) Simulate the optimal auction OPT_F on the first n bidders $1, 2, \dots, n$.
- (2) If the item was not awarded in step 1, then give the item to bidder $n+1$ for free.

We defined \mathcal{A} to possess two properties useful for the analysis. First, its expected revenue (with $n+1$ bidders) is exactly that of OPT_F (with n bidders). Second, \mathcal{A} always allocates the item.

We can finish the proof by arguing that, for i.i.d. regular bidder valuations, the Vickrey auction maximizes expected revenue over all auctions that are guaranteed to allocate the item. By the equivalence of expected revenue and expected virtual welfare, the optimal auction that always allocates the item awards the item to the bidder with the highest virtual valuation (even if this is negative). The Vickrey auction awards the item to the bidder with the highest valuation. Since bidders’ valuations are i.i.d. draws from a regular distribution, all bidders share the same increasing virtual valuation function φ . Thus the bidder with the highest virtual valuation is always the bidder with the highest valuation. We conclude that 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_F (with n bidders). ■

³That is, OPT_F is the Vickrey auction with the monopoly reserve price $\varphi^{-1}(0)$, where φ is the virtual valuation function of F .

5 Case Study: Reserve Prices in Yahoo! Keyword Auctions

So how does all this optimal auction theory get used, anyway? We next discuss a 2008 field experiment by Ostovsky and Schwarz [6], which explored whether or not the lessons of auction theory could be used to increase revenue for Yahoo! in its keyword search auctions.

Recall from Lecture #2 the standard model of keyword auctions. Which such auction maximizes the expected revenue, at least in theory? Assuming that bidders' valuations-per-click are drawn i.i.d. from a regular distribution F , it is simply to rank bidders by bid (from the best slot to the worst) after applying the monopoly reserve price $\varphi^{-1}(0)$ to all bidders, where φ is the virtual valuation function of F . See the exercises for details.

What had Yahoo! been doing, up to 2008? First, they were using relatively low reserve prices — initially \$.01, later \$.05, and \$.10 in 2008. Perhaps more naively, they were using the same reserve price of \$.10 across all keywords, even though some keywords surely warranted higher reserve prices than others (e.g., compare the searches “divorce lawyer” with “pizza”). How would Yahoo!’s revenue change if reserve prices were changed, independently for each keyword, to be theoretically optimal?

The field experiment described in [6] had two parts. First, a lognormal valuation distribution was posited for each of a half million keywords based on past bidding data.⁴ This step is somewhat ad hoc but there is no evidence that the final conclusions depend on its details (such as the particular family of distributions used).

Next, theoretically optimal reserve prices were computed for each keyword, assuming valuations are drawn from the fitted distributions. As expected, the optimal reserve price varies a lot across keywords, but there are plenty of keywords with a theoretically optimal reserve price of \$.30 or \$.40. Thus, Yahoo!’s uniform reserve price was much too low, relative to the theoretical advice, on many keywords.

The obvious experiment is to try out the theoretically optimal (and generally higher) reserve prices to see how they do. Yahoo!’s top brass wanted to be a little more conservative, though, and set the new reserve prices to be the average of the old ones (\$.10) and the theoretically optimal ones.⁵ And the change worked: auction revenues went up several per cent (of a very large number). The new reserve prices were especially effective in markets that are valuable but “thin,” meaning not very competitive (less than 6 bidders). Better reserve prices were credited by Yahoo!’s president as the biggest reason for higher search revenue in Yahoo!’s third-quarter report in 2008.

⁴Since Yahoo!, like other search engines, uses a non-DSIC auction based on the GSP auction (see Problem 3), one cannot expect the bids to be truthful. In [6], valuations are reversed engineered from the bids under the assumption that bidders are playing the equilibrium that is outcome-equivalent to the dominant-strategy outcome of the DSIC auction (as in Problem 3).

⁵It turns out that, both in theory and empirically, this initial change accounts for most of the revenue increase. Increasing the reserve price further does not have much effect on revenue.

References

- [1] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.
- [2] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, pages 129–138, 2010.
- [3] J. D. Hartline. Mechanism design and approximation. Book draft. October, 2013.
- [4] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 225–234, 2009.
- [5] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [6] M. Ostrovsky and M. Schwarz. Reserve prices in internet advertising auctions: A field experiment. Working paper, December 2009.
- [7] E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability*, 12(4):1213–1216, 1984.