# CS269I: Incentives in Computer Science
# Lecture #16: Revenue-Maximizing Auctions*

Tim Roughgarden†

November 16, 2016

## 1 Revenue Maximization and Bayesian Analysis

Thus far, we've focused on the design of auctions that maximize the social welfare (at least in the case where bids are truthful). Welfare maximization is a first-order concern in many scenarios, and as we've seen, this objective has guided the design of auctions for sponsored search and online advertising more generally. Revenue is generated in welfare-maximizing auctions, but only as a side effect, a necessary evil to incentivize participants to report their private information. (Remember what happened when we tried to charge nothing in a single-item auction?)

What if our goal is to maximize the seller's revenue, rather than the social welfare? This lecture, our last on auction design, explains the changes that need to be made.

### 1.1 One Bidder and One Item

The following trivial example is illuminating. Suppose there is one item and only one bidder, with a private valuation $v$. With only one bidder, the space of truthful auctions is small: they are precisely the "posted prices," or take-it-or-leave-it offers.[1] If the seller posts a price of $r$, then her revenue is either $r$ (if $v \geq r$) or 0 (if $v < r$).

Maximizing the welfare in this setting is trivial: just set $r = 0$, so that you always give the item to the bidder for free. (This can be thought of as a Vickrey auction with one bidder.) Note that this welfare-maximizing posted price is *independent of* $v$.

Suppose we wanted to maximize revenue. How should we set $r$? If we telepathically knew $v$, then we would set $r = v$. But with $v$ private, what should we do? It's not obvious how to reason about this question.

---



†Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: `tim@cs.stanford.edu`.

[1]One can also randomize over posted prices, but the point of the example remains the same.

The fundamental issue is that, for the revenue objective, different auctions do better on different inputs. Here's another way to think about it: if your goal is to maximize social welfare, you'll never regret setting $r$ to 0, even after you find out what the bidder would have been willing to pay. If your goal is to maximize revenue, then unless you are very lucky, in hindsight you'll wish you had set a different value of $r$ (namely $r = v$).

## 1.2 Bayesian Analysis

Comparing different auctions for revenue maximization requires a model to reason about trade-offs across different inputs. The standard model for doing this in economics is *average-case* or *Bayesian* analysis. For now, our model comprises the following ingredients:

- A single-item auction (1 seller, 1 item, $n$ bidders).
- The private valuations $v_1, \ldots, v_n$ of the bidders are assumed to be independently and identical distributed (i.i.d.) according to a distribution $F$.[2]
- The distribution $F$ is known in advance to the seller or auction designer. For example, $F$ could be derived from past data (see also Section 3). The realizations $v_1, \ldots, v_n$ of bidders' valuations are private, as usual.

The setup can be generalized in various ways. First, we'll see how to go beyond single-item auctions later, in the context of sponsored search. Second, the theory we'll describe remains well understood when bidders' valuations are drawn independently from (possibly different) distributions $F_1, \ldots, F_n$. Today we'll stick with the i.i.d. case because it is simpler; see the instructor's CS364A lecture notes (or [5]) for a treatment of the non-identical case. Third, the theory becomes significantly more complicated when bidders' valuation can be correlated. One often makes the assumption of independent valuations not because one literally believes it, but rather because the assumption results in a helpful theoretical framework for thinking about auction design. Finally, in this lecture we consider only truthful auctions. Note that this implies that bidders don't need to know or care about the valuation distribution $F$ (they should all just bid truthfully in any case). Again, the theory is more general, and actually implies that there is no loss of generality in restricting attention to truthful auctions. For example, you can't make more money with a first-price auction (at equilibrium) then with a second-price auction.[3]

In a Bayesian environment, it is clear how to define the "optimal" auction — it is the auction that, among all truthful auctions, has the highest expected revenue, where the expectation is with respect to the randomness in bidders' valuations.

[2]Recall $F(z)$ denotes the probability that a random variable with distribution $F$ has value at most $z$.

[3]Cf., the discussion of revenue equivalence in Lecture #14.

## 1.3 One Bidder and One Item, Revisited

With our Bayesian model, single-bidder single-item auctions are now easy to reason about. The expected revenue of a posted price $r$ is simply

$$\underbrace{r}_{\text{revenue of a sale}} \cdot \underbrace{(1 - F(r))}_{\text{probability of a sale}}.$$

Given a distribution $F$, it is usually a simple matter to solve for the best $r$. The optimal posted price is called the *monopoly price* of the distribution $F$. Since truthful auctions are posted prices (and randomizations thereof), posting the monopoly price is the revenue-maximizing auction. For instance, if $F$ is the uniform distribution on $[0, 1]$ (i.e., $F(x) = x$ on $[0, 1]$), then setting the derivative of the expected revenue to zero and solving shows that the monopoly price is $\frac{1}{2}$ (with an expected revenue of $\frac{1}{4}$).

The plot thickens even with two bidders, where the space of truthful auctions is larger. For example, consider a single-item auction with two bidders with valuations drawn i.i.d. from the uniform distribution on $[0, 1]$. We could of course run the Vickrey auction; its revenue is the expected value of the smaller bid, which is $\frac{1}{3}$ (see Lecture #14).

We could also supplement the Vickrey auction with a *reserve price*, analogous to the "opening bid" in an eBay auction. In a Vickrey auction with reserve $r$, the winner is the highest bidder, unless all bids are less than $r$, in which case there is no winner. The winner (if any) pays the second-highest bid or $r$, whichever is larger. (Check that this is a truthful auction.) From a revenue standpoint, adding a reserve price $r$ is both good and bad: you lose revenue when all bids are less than $r$, but you gain revenue when exactly one bid is above $r$ (since the selling price is higher). In our case, adding a reserve price of $\frac{1}{2}$ turns out to be a net gain, raising the expected revenue from $\frac{1}{3}$ to $\frac{5}{12}$ (exercise). But can we do better? Perhaps using a different reserve price?[4] Perhaps usually a totally different auction format? While the rich space of truthful auctions makes this an intimidating question, the rest of this lecture provides a complete solution, originally given by Myerson [3].

# 2 Optimal Auction Theory

We need one technical definition to formally state our main result.[5]

**Definition 2.1** A valuation distribution $F$ with density $f$ is *regular* if

$$x - \frac{1 - F(x)}{f(x)} \tag{1}$$

is a nondecreasing function of $x$.[6]

[4] For example, in a first-price auction, the more competition there is, the more aggressively you should bid (Lecture #14). So maybe the more competition there is, the higher a reserve price the seller should use?

[5] This definition is for completeness and accuracy, and it's not important that you remember it precisely. It is good to remember some examples and non-examples, however.

[6] In auction theory, the expression (1) is called the *virtual valuation* function of the distribution $F$.

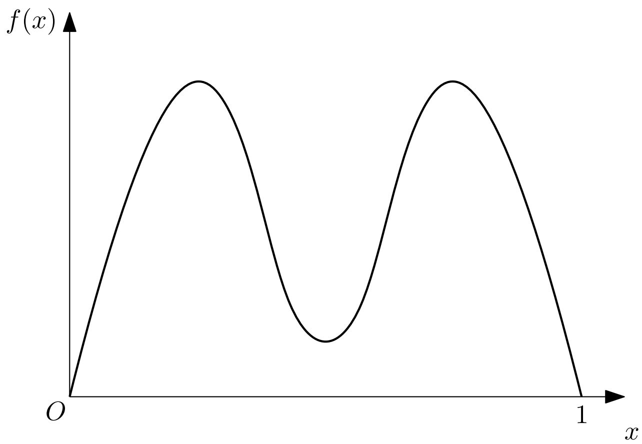


Figure 1: A bimodal distribution is not regular in general.

For example, the uniform distribution on $[0,1]$ satisfies $F(x) = x$ and $f(x) = 1$, so the expression (1) reads $2x - 1$ (for $x \in [0,1]$). This is increasing in $x$, so this distribution is regular. Other examples include exponential distributions, lognormal distributions, and power-law distributions with exponent bigger than 2. The most natural example of a distribution that is *not* regular is a multi-modal distribution with distinctly different modes (Figure 1). Do you see why such a distribution is not regular in general?

For the rest of this lecture, we restrict attention to regular valuation distributions. The theory of revenue-maximizing auctions has been thoroughly worked out also in the case of general (non-regular) distributions, but the main results are simpler to state for the common case of regular distributions.

The following is a major result.[7]

**Theorem 2.2 ([3])** *If bidders' valuations are drawn i.i.d. from a regular distribution $F$, then the expected revenue-maximizing auction is the Vickrey auction, with a reserve price equal to the monopoly price of $F$ (i.e.* $\operatorname{argmax}_r r(1 - F(r))$*).*

Theorem 2.2 is really just a special case of the general theory in [3] and can be extended to non-identical and non-regular valuation distributions. In the interests of time, we won't prove Theorem 2.2 here. (It takes about a lecture; see the instructor's CS364A lecture notes or [5].)

Theorem 2.2 implies that, for i.i.d. bidders and a regular valuation distribution, eBay (with a suitable opening bid) is the optimal auction format! Given the richness of the auction design space, it is amazing that such a simple and practical auction pops out as the optimal one. For instance, in the example with two bidders with uniformly distributed valuations (Section 1.3), the expected revenue of $\frac{5}{12}$ achieved by the Vickrey auction with reserve price $\frac{1}{2}$ really is the maximum possible (since $\frac{1}{2}$ is the monopoly price for the uniform distribution).

[7] Myerson received the 2007 Nobel Prize in Economics, in part for this work. To this day, the theory of revenue-maximizing auctions is called "optimal auction theory," following the title of [3].

Another remarkable aspect of Theorem 2.2 is that the optimal reserve price is *independent of* $n$, the number of bidders. (So the analogy suggested in Section 1.3 with bidding in a first-price auction is incorrect.) The reserve price is less and less likely to be relevant as $n$ grows, however, since most of the time there will be at least two bidders who clear it. (See also Section 4.)

Theorem 2.2 extends beyond single-item auctions. For example, suppose there are $k$ identical copies of an item, and each bidder wants at most one copy (as in Exercise 29). Recall that the social welfare-maximizing auction gives a copy of the item to each of the top $k$ bidders, at a price equal to the highest losing bid (the $(k+1)$th highest bid overall). With valuations drawn i.i.d. from a regular distribution $F$, the revenue-maximizing auction differs only in the addition of a reserve price $r$, again equal to the monopoly price of $F$. (So the optimal reserve price is now not only independent of the number of bidders, but also of the number of items $k$.) Formally, the winners in the latter auction are the bidders that are both in the top $k$ bidders and bid above the reserve $r$, and the selling price of each copy is the maximum of $r$ and the $(k+1)$th highest bid. (Check that this auction is truthful.)

The revenue-maximizing sponsored search auction, with bidders' valuations-per-click drawn i.i.d. from a regular distribution, follows the same template: just add a reserve price equal to the monopoly price to the welfare-maximizing (VCG) auction. Formally, the winners are the bidders that are both in the top $k$ bidders and bid above the reserve $r$, with the $i$th highest winning bidder assigned to the $i$ slot, and the price-per-click of each slot is the maximum of $r$ and what the VCG price would have been for that slot. Again, this is a truthful auction, and the reserve price depends only on $F$, and not on any of the other details of the setting (like the click-through rates of different slots).

# 3 Case Study: Reserve Prices in Yahoo Sponsored Search Auctions

So how does all this optimal auction theory get used, anyway? We next discuss a 2008 field experiment by Ostrovsky and Schwarz [4], which explored whether or not the lessons of auction theory could be used to increase revenue for Yahoo in its sponsored search auctions. As noted above, assuming that bidders' valuations-per-click are drawn i.i.d. from a regular distribution $F$, the revenue-maximizing auction simply ranks bidders by bid (from the best slot to the worst) after applying the monopoly reserve price to all bidders.

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 [4] had two parts. First, a lognormal valuation distribution was posited for each of roughly a half million keywords based on past bidding data.[8] 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 varied 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.[9] And the change worked: auction revenues went up several percent (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.

# 4 Prior-Independent Auctions and the Bulow-Klemperer Theorem

Theorem 2.2 identifies the revenue-maximizing auction as a function of the valuation distribution $F$. (Different $F$'s have different monopoly prices and hence different optimal auctions.) But how do we know $F$? What if we don't have very much recent and relevant data about bidders' valuations? Can we do almost as well as if we did? The goal in *prior-independent* auction design is to design an auction, whose description is independent of the underlying distribution, that performs almost as well as if the distribution were known in advance [2]. This has been a hot topic on the boundary of computer science and economics throughout this decade.

Today, we'll only have time to mention one beautiful result from classical auction theory which is also an important precursor to the theory 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.

---

[8]Since Yahoo, like other search engines, uses the non-truthful GSP auction, one cannot expect the bids to be truthful. In [4], valuations were reversed engineered from the bids under the assumption that bidders are playing the equilibrium that is outcome-equivalent to the truthful outcome of the VCG auction (Lecture #15).

[9]Both in theory and empirically, this more conservative change accounts for most of the revenue increase. There are usually diminishing returns to revenue as the reserve price approaches the theoretical optimum, providing flexibility near the optimal price. The intuition for this principle is that the derivative of the expected revenue with respect to the reserve price is 0 at the optimal point.

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

$$\mathbf{E}_{v_1,\ldots,v_{n+1}\sim F}[\mathit{Rev(VA)}\ (n+1\ \mathit{bidders})] \geq \mathbf{E}_{v_1,\ldots,v_n\sim F}[\mathit{Rev(OPT_F)}\ (n\ \mathit{bidders})], \tag{2}$$

*where VA and $OPT_F$ denote the Vickrey single-item auction (with no reserve price) and the optimal single-item auction for $F$, respectively.*[10]

Notice that the auction in the left-hand side of (2) — 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 (2) 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 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 additional serious participants, rather than sharpening your knowledge of bidders' preferences (of course, do both if you can!).

We'll skip the proof of Theorem 4.1; see the instructor's CS364A lecture notes or [5] for details. (Once Theorem 2.2 has been proved, the proof of Theorem 4.1 is pretty short.)

The guarantee in Theorem 4.1 also implies that, as the number of i.i.d. bidders grows in a single-item auction, the Vickrey auction has expected revenue almost as high as the optimal auction (with the same number of bidders). One intuition for this result is that the reserve price in the optimal auction—the only difference between the two auctions—is ever-less likely to kick in as $n \to \infty$.

**Corollary 4.2** *In a single-item auction with $n \geq 2$ bidders with valuations drawn i.i.d. from a regular distribution $F$, the expected revenue of the Vickrey auction is at least $\frac{n-1}{n}$ times that of the optimal auction of $F$.*

Corollary 4.2 follows immediately from chaining together the Bulow-Klemperer theorem (Theorem 4.1) and the following lemma.

**Lemma 4.3** *For every regular valuation distribution $F$ and $n \geq 1$,*

$$\mathbf{E}_{v_1,\ldots,v_{n-1}\sim F}[\mathit{Rev(OPT_F)}\ (n-1\ \mathit{bidders})] \geq \frac{n-1}{n}\mathbf{E}_{v_1,\ldots,v_n\sim F}[\mathit{Rev(OPT_F)}\ (n\ \mathit{bidders})]. \tag{3}$$

*Proof:* We use a simulation argument. Fix $F$ and $n \geq 1$. Let $A^*$ denote the optimal truthful auction for $n$ bidders. (The Vickrey auction with reserve price equal to the monopoly price of $F$.) We proceed to define an auction $\hat{A}$ on $n-1$ bidders, just for the sake of the proof. We want $\hat{A}$ to simulate $A^*$, which means that $\hat{A}$ needs to "make up" an extra bidder for $A^*$'s sake. Formally, here is the description of $\hat{A}$:

[10] That is, $OPT_F$ is the Vickrey auction with a reserve price equal to the monopoly price of $F$.

1. Accept bids $b_1, b_2, \ldots, b_{n-1}$ from the $n-1$ bidders.

2. Define the fictitious bid $b_n$ as a random sample drawn from $F$.

3. Run $A^*$ on the bid vector $b_1, \ldots, b_n$.

4. If $A^*$ sells the item to bidder $i \in \{1, 2, \ldots, n-1\}$ at the price $p$, or to no bidder, then $\hat{A}$ copies the outcome of $A^*$.

5. If $A^*$ sells the item to the $n$th bidder, then $\hat{A}$ does not sell the item to anyone.

As you should check, the auction $\hat{A}$ inherits truthfulness from $A^*$.

By symmetry, the auction $A^*$ receives an equal fraction of its expected revenue from each of the $n$ bidders. (Each bidder is equally likely to be the winner, and the distribution of the selling price conditioned on winning is the same for every bidder.) The auction $\hat{A}$ gets the same expected revenue as $A^*$ from each of the first $n-1$ bidders, but none of $A^*$'s revenue from the $n$th bidder. This means that the expected revenue of $\hat{A}$ is exactly $\frac{n-1}{n}$ times that of $A^*$. The optimal auction for $n-1$ bidders can only be better. ■

Approximation guarantees for simple and practical auctions (like Corollary 4.2) have been a hot topic in computer science for a number of years, and are also becoming more prevalent in economics. Such results are particularly important in more complex settings where the optimal auction is not as simple as in Theorem 2.2 (like with non-i.i.d. bidders, or with multiple non-identical items).

# 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. *Games and Economic Behavior*, 91:318–333, 2015.

[3] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[4] M. Ostrovsky and M. Schwarz. Reserve prices in Internet advertising auctions: A field experiment. Working paper, December 2009.

[5] T. Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016.