

CS364A: Algorithmic Game Theory

Lecture #5: Revenue-Maximizing Auctions*

Tim Roughgarden[†]

October 7, 2013

1 The Challenge of Revenue Maximization

1.1 Welfare-Maximization, Revisited

Thus far, we’ve focused on the design of mechanisms that maximize, exactly or approximately, the *welfare* objective

$$\sum_{i=1}^n v_i x_i \tag{1}$$

over all feasible outcomes (x_1, \dots, x_n) in some set X . Revenue is generated in welfare-maximizing auctions, but only as a side effect, a necessary evil to incentivize participants to report their private information. This lecture begins our discussion of auctions that are explicitly designed to raise as much revenue as possible.

We started with the welfare objective for several reasons. One is that it’s a fundamental objective function, relevant to many real-world scenarios. For instance, in government auctions (e.g., to sell wireless spectrum), the primary objective is welfare maximization — revenue is also important but is usually not the first-order objective. Also, in competitive markets, it is often thought that a seller should focus on welfare-maximization, since otherwise someone else will (potentially stealing their customers).

The second reason we started with welfare-maximization is pedagogical: welfare is special. In every single-parameter environment (and even more generally, see Lecture #7), there is a DSIC mechanism for maximizing welfare *ex post* — as well as if the designer knew all of the private information (the v_i ’s) in advance. In this sense, one can satisfy the DSIC constraint “for free.” This is an amazingly strong performance guarantee, and it cannot generally be achieved for other objective functions.

*©2013, Tim Roughgarden.

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

1.2 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 direct-revelation DSIC 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 its 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. Note that this optimal posted price is *independent of v* .

Suppose we wanted to maximize revenue. How should we set r ? Ex post (i.e., telepathically knowing v), we would set $r = v$. But with v private, what should we do? It’s not obvious how to reason about this question.

The fundamental issue is that, for the revenue objective, different auctions do better on different inputs. With a single item and bidder, a posted price of 20 will do very well on inputs where v is 20 or a little larger, and terribly on smaller inputs (for which smaller posted prices will do better). Such trade-offs are familiar to students of algorithms. For example, different sorting algorithms (e.g., InsertionSort vs. QuickSort) run faster on different inputs. Different heuristics for the Traveling Salesman Problem (e.g., local search vs. linear programming) have smaller solution error on different inputs. Welfare-maximization, where there is an input-independent optimal DSIC mechanism, is special indeed.

1.3 Bayesian Analysis

Comparing different auctions for revenue maximization requires a model to reason about trade-offs across different inputs. Today we introduce the most classical and well-studied model for doing this: *average-case* or *Bayesian* analysis. Our model comprises the following ingredients:

- A single-parameter environment (see Lecture #3).
- The private valuation v_i of participant i is assumed to be drawn from a distribution F_i with density function f_i with support contained in $[0, v_{\max}]$.² We assume that the distributions F_1, \dots, F_n are independent (but not necessarily identical). In practice, these distributions are typically derived from data, such as bids in past auctions.
- The distributions F_1, \dots, F_n are known in advance to the mechanism designer. The realizations v_1, \dots, v_n of bidders’ valuations are private, as usual. Since we focus on DSIC auctions, where bidders have dominant strategies, the bidders do not need to know the distributions F_1, \dots, F_n .³

¹Precisely, these are the deterministic such auctions. One can also randomize over posted prices, but the point of the example remains the same.

²Recall $F_i(z)$ denotes the probability that a random variable drawn from F_i has value at most z .

³The main results in today’s lecture apply more generally to “Bayes-Nash incentive-compatible” auctions; in this case, the bidders must also know the distributions F_1, \dots, F_n .

In a Bayesian environment, it is clear how to define the “optimal” auction — it is the auction that, among all DSIC auctions, has the highest expected revenue, where the expectation is with respect to the given distribution $F_1 \times F_2 \times \cdots \times F_n$ over valuation profiles \mathbf{v} (assuming truthful bidding).

1.4 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 DSIC mechanisms 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 the monopoly price is $\frac{1}{2}$, achieving an expected revenue of $\frac{1}{4}$.

The plot thickens even with two bidders, where the space of DSIC 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 $\frac{1}{3}$ (exercise).

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 allocation rule awards the item to the highest bidder, unless all bids are less than r , in which case no one gets the item. The corresponding payment rule charges the winner (if any) the second-highest bid or r , whichever is larger. From a revenue vantagepoint, 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, or perhaps usually a totally different auction format? While the rich space of DSIC auctions makes this an intimidating question, the rest of this lecture provides a complete solution, originally given by Myerson [2].

2 Expected Revenue Equals Expected Virtual Welfare

Our goal is to characterize the optimal (i.e., expected revenue-maximizing) DSIC auction for every single-parameter environment and distributions F_1, \dots, F_n .⁴ We begin with a preliminary observation.

⁴We won’t prove it today, but the auctions we identify are optimal in a much stronger sense; see the discussion in Section 3.2.

Step 0: By the Revelation Principle from last lecture, every DSIC auction is equivalent to — and hence has the same expected revenue as — a direct-revelation DSIC mechanism (\mathbf{x}, \mathbf{p}) . We can therefore consider only direct-revelation mechanisms from here on. We correspondingly assume truthful bids (i.e., $\mathbf{b} = \mathbf{v}$) for the rest of the lecture.

The expected revenue of an auction (\mathbf{x}, \mathbf{p}) is

$$\mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \mathbf{p}(\mathbf{v}) \right], \quad (2)$$

where the expectation is with respect to the distribution $F_1 \times \dots \times F_n$ over bidders' valuations. It is not clear how to directly maximize the expression (2) over the space of DSIC mechanisms. In this section, we derive a *second* formula for the expected revenue of an auction. This second formula only references the allocation rule of a mechanism, and not its payment rule, and for this reason has a form that is far easier to maximize.

As a starting point, recall Myerson's payment formula from Lecture #3:

$$p_i(b_i, \mathbf{b}_{-i}) = \int_0^{b_i} z \cdot x'_i(z, \mathbf{b}_{-i}) dz. \quad (3)$$

We derived this equation assuming that the allocation function $x_i(z, \mathbf{b}_{-i})$ is differentiable. By standard advanced calculus, the same formula holds more generally for an arbitrary monotone function $x_i(z, \mathbf{b}_{-i})$, including piecewise constant functions, for a suitable interpretation of the derivative $x'_i(z, \mathbf{b}_{-i})$ and the corresponding integral. Similarly, all of the following proof steps, which make use of calculus maneuvers like integration by parts, can be made fully rigorous for arbitrary bounded monotone functions without significant difficulty. We leave the details to the interested reader.⁵

Equation (3) states that payments are fully dictated by the allocation rule. Thus, at least in principle, we can express the expected revenue of an auction purely in terms of its allocation rule, with no explicit reference to its payment rule. Will the resulting revenue formula will be easier to maximize than the original one? It's hard to know without actually doing it, so let's do it.

Step 1: Fix i and \mathbf{v}_{-i} ; recall that \mathbf{v}_{-i} is a random variable (as is v_i), and we'll integrate out over it later.

By Myerson's payment formula (3), we can write the expected payment by bidder i for a given value of \mathbf{v}_{-i} as

$$\mathbf{E}_{v_i \sim F_i} [p_i(\mathbf{v})] = \int_0^{v_{\max}} p_i(\mathbf{v}) f_i(v_i) dv_i = \int_0^{v_{\max}} \left[\int_0^{v_i} z \cdot x'_i(z, \mathbf{v}_{-i}) dz \right] f_i(v_i) dv_i.$$

Note that in the first equality we're exploiting the independence of bidders' valuations — the fixed value of \mathbf{v}_{-i} has no bearing on the distribution F_i from which v_i is drawn.

⁵For example, every bounded monotone function is integrable, and is differentiable except on a set of measure zero.

This step is exactly what we knew was possible in principle — rewriting the payment in terms of the allocation rule. For this to be useful, we need some simplifications.

Step 2: Whenever you have a double integral (or double sum) that you don't know how to interpret, it's worth reversing the integration order. Here, reversing the order of integration leads to a nice simplification, suggesting we're on the right track:

$$\begin{aligned} \int_0^{v_{\max}} \left[\int_0^{v_i} z \cdot x'_i(z, \mathbf{v}_{-i}) dz \right] f_i(v_i) dv_i &= \int_0^{v_{\max}} \left[\int_z^{v_{\max}} f_i(v_i) dv_i \right] z \cdot x'_i(z, \mathbf{v}_{-i}) dz \\ &= \int_0^{v_{\max}} (1 - F_i(z)) \cdot z \cdot x'_i(z, \mathbf{v}_{-i}) dz. \end{aligned}$$

Step 3: Integration by parts is also worth trying when attempting to massage an integral into a more interpretable form, especially if there's an obvious derivative hiding in the integrand. Here, we again get some encouraging simplifications:

$$\begin{aligned} &\int_0^{v_{\max}} \underbrace{(1 - F_i(z)) \cdot z}_f \cdot \underbrace{x'_i(z, \mathbf{v}_{-i})}_{g'} dz \\ &= \underbrace{(1 - F_i(z)) \cdot z \cdot x_i(z, \mathbf{v}_{-i}) \Big|_0^{v_{\max}}}_{=0-0} - \int_0^{v_{\max}} x_i(z, \mathbf{v}_{-i}) \cdot (1 - F_i(z) - z f_i(z)) dz \\ &= \int_0^{v_{\max}} \underbrace{\left(z - \frac{1 - F_i(z)}{f_i(z)} \right)}_{:=\varphi_i(z)} x_i(z, \mathbf{v}_{-i}) f_i(z) dz. \end{aligned}$$

Notice that we can interpret the final expression as an expected value, where z is drawn from the distribution F_i .

Step 4: To simplify and help interpret the expression above, we introduce some new notation. The *virtual valuation* $\varphi_i(v_i)$ of bidder i with valuation v_i drawn from F_i is

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

Note that the virtual valuation of a bidder depends on its own valuation and distribution, and not on those of the others.

For example, consider a bidder i with valuation drawn from the uniform distribution on $[0, 1]$. Then $F_i(z) = z$, $f_i(z) = 1$, and $\varphi_i(z) = z - \frac{1-z}{1} = 2z - 1$ in $[0, 1]$. Notice that a virtual valuation can be negative! See the exercises for more examples.

Steps 1–4 Summary: For every bidder i and valuations \mathbf{v}_{-i} ,

$$\mathbf{E}_{v_i \sim F_i} [p_i(\mathbf{v})] = \mathbf{E}_{v_i \sim F_i} [\varphi_i(v_i) \cdot x_i(\mathbf{v})]. \quad (4)$$

Remark 2.1 Virtual valuations play a central role in the design of Bayesian optimal auctions. Is there any intuition for what they mean? One coarse way to interpret the formula

$$\varphi_i(v_i) = \underbrace{v_i}_{\text{what you'd like to charge } i} - \underbrace{\frac{1 - F_i(v_i)}{f_i(v_i)}}_{\text{"information rent" earned by bidder } i}$$

is to think of v_i as the maximum revenue obtainable from bidder i , and the second term as the inevitable revenue loss caused by not knowing v_i in advance (a.k.a. “information rent”). A second and more accurate interpretation of $\varphi_i(v_i)$ is as the slope of a “revenue curve” at v_i , where the revenue curve plots the expected revenue obtained from an agent with valuation drawn from F_i , as a function of the probability of a sale. The exercises elaborate on this second interpretation.

Step 5: Take the expectation, with respect to \mathbf{v}_{-i} , of both sides of (4) to obtain:

$$\mathbf{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\varphi_i(v_i) \cdot x_i(\mathbf{v})].$$

Step 6: Apply linearity of expectations (twice) to finish the derivation:

$$\mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n p_i(\mathbf{v}) \right] = \sum_{i=1}^n \mathbf{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \sum_{i=1}^n \mathbf{E}_{\mathbf{v}}[\varphi_i(v_i) \cdot x_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \varphi_i(v_i) \cdot x_i(\mathbf{v}) \right]. \quad (5)$$

The final term in (5) is our second formula for the expected revenue of an auction, and we should be pleased with its relative simplicity. Note that if we removed the φ_i ’s from the expression, we would be left with an old friend: the expected *welfare* of the auction. For this reason, we refer to $\sum_{i=1}^n \varphi_i(v_i) \cdot x_i(\mathbf{v})$ as the *virtual welfare* of an auction on the valuation profile \mathbf{v} . We have proved that, for every auction,

$$\text{EXPECTED REVENUE} = \text{EXPECTED VIRTUAL WELFARE}. \quad (6)$$

In particular, maximizing expected revenue over the space of DSIC auctions reduces to maximizing expected virtual welfare.

3 Bayesian Optimal Auctions

It is shocking that a formula as simple as (6) holds. It says that even though we only care about payments, we can focus on an optimization problem that concerns only the mechanism’s allocation rule. This second form is far more operational, and we proceed to determine the auctions that maximize it.

3.1 Maximizing Expected Virtual Welfare

As a warm up, let's make two extra assumptions. First, consider a single-item auction. Second, assume that the bidders are i.i.d. That is, all F_i 's are a common F , and thus all virtual valuation functions φ_i are the same.

How should we choose the allocation rule \mathbf{x} to maximize the expected virtual welfare

$$\mathbf{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_{i=1}^n \varphi_i(v_i) x_i(\mathbf{v}) \right] ? \quad (7)$$

We have the freedom of choosing $\mathbf{x}(\mathbf{v})$ for each input \mathbf{v} , and have no control over the input distribution \mathbf{F} or the virtual values $\varphi_i(v_i)$. Thus, the obvious approach is to maximize pointwise: separately for each input \mathbf{v} , we choose $\mathbf{x}(\mathbf{v})$ to maximize the virtual welfare $\sum_{i=1}^n \varphi_i(v_i) x_i(\mathbf{v})$ obtained on the input \mathbf{v} (subject to feasibility of $(x_1, \dots, x_n) \in X$). We call this the *virtual welfare-maximizing allocation rule*.

In a single-item auction, where the feasibility constraint is $\sum_{i=1}^n x_i(\mathbf{v}) \leq 1$ for each \mathbf{v} , the virtual welfare-maximizing rule just awards the item to the bidder with the highest virtual valuation. Well not quite: recall that virtual valuations can be negative (e.g., $\varphi_i(v_i) = 2v_i - 1$ when v_i is uniform between 0 and 1), and if every bidder has a negative virtual valuation then the virtual welfare is maximized by not awarding the item to anyone. (We already saw in the single-bidder example that maximizing revenue entails not selling the item in some cases.)

Choosing $\mathbf{x}(\mathbf{v})$ separately for each \mathbf{v} to maximize $\sum_{i=1}^n \varphi_i(v_i) x_i(\mathbf{v})$ defines an allocation rule that maximizes the expected virtual welfare (7) over all allocation rules (monotone or not). The key question is: *is this virtual welfare-maximizing rule monotone?* If so, then it can be extended to a DSIC auction, and by (6) this auction has the maximum-possible expected revenue.

The answer to this key question depends on the valuation distribution F . If the corresponding virtual valuation function φ is increasing, then the virtual welfare-maximizing allocation rule is monotone.

Definition 3.1 A distribution F is *regular* if the corresponding virtual valuation function $v - \frac{1-F(v)}{f(v)}$ is strictly increasing.

For most applications, Definition 3.1 can be relaxed to allow nondecreasing virtual valuation functions.

We saw that the uniform distribution on $[0, 1]$ has virtual valuation function $2v - 1$ and hence is regular. So are other uniform distributions, exponential distributions, and lognormal distributions. Irregular distributions include many multi-modal distributions and distributions with sufficiently heavy tails. See the exercises for concrete examples.

Let's return to a single-item auction with i.i.d. bidders, under the additional assumption that the valuation distribution is regular. The virtual-welfare maximizing allocation rule, which allocates to the bidder with highest nonnegative virtual valuation (if any), is monotone and yields the optimal auction. Moreover, since all bidders share the same increasing virtual

valuation function, the bidder with the highest virtual valuation is also the bidder with the highest valuation. This allocation rule is thus equivalent to the Vickrey auction with a reserve price of $\varphi^{-1}(0)$. Thus, 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 DSIC auction design space, it is amazing that such a simple and practical auction pops out as the optimal one.

More generally, consider an arbitrary single-parameter environment and valuation distributions F_1, \dots, F_n . The virtual welfare-maximizing allocation rule is now defined as that which, for each input \mathbf{v} , chooses the feasible allocation that maximizes the virtual welfare $\sum_{i=1}^n \varphi_i(v_i) x_i(\mathbf{v})$. If every distribution F_i is regular, then this allocation rule is monotone (see the exercises). Coupling it with the unique payment rule to meet the DSIC constraint, we obtain the optimal auction. In this sense, we have solved the Bayesian optimal auction problem for every single-parameter environment with regular valuation distributions.

3.2 Extensions

The theory developed in this lecture, which is due to Myerson [2], is even more general. First, it can be extended to accommodate valuation distributions that are not regular.⁶ Since the virtual welfare-maximization allocation rule is not monotone in this case, one has to work harder and solve for the monotone allocation rule with the maximum expected virtual welfare. This can be done by “ironing” virtual valuation functions to make them monotone, while at the same time preserving the virtual welfare of the auctions that matter. See Hartline [1, Chapter 3] for a textbook treatment.

Second, while today’s lecture restricted attention to DSIC auctions for simplicity, the (DSIC) optimal auctions we identified are optimal even amongst the much larger set of “Bayesian incentive compatible” mechanisms. For example, first-price auction formats cannot achieve more revenue (at equilibrium) than the best DSIC auction. Thus, for revenue-maximization in single-parameter problems, the DSIC constraint comes for free. This extension does not require significant new ideas beyond what we covered today, and we’ll discuss it in CS364B.

References

- [1] J. D. Hartline. Mechanism design and approximation. Book draft. October, 2013.
- [2] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

⁶Independence of the distributions, however, is crucial and cannot be relaxed without significantly changing the results.