Game Theory and Algorithms*
Lecture 13: Maximizing Auctioneer Revenue

April 14, 2011

Summary: We introduce seminal results of Myerson [2] on "optimal mechanism design." This pertains to the case that each player in the auction is drawn from a random distribution which is known in advance. Our goal is to maximize the revenue of the auctioneer, while keeping the auction truthful-in-expectation. One very interesting consequence is that for a broad class of distributions, the *income-maximizing* auction for n bidders from this distribution is just the second-price auction with a reserve price which is independent of n.

1 Single-Parameter Auctions

We will consider today "single-parameter auctions." This means that for each bidder, there are only two kinds of outcomes as far as they are concerned, winning ones and losing ones. Hence, each alternative is equivalent to specifying a set of "winners" and the rest of the players as losers. Without loss of generality, we assume players have 0 utility for losing; they have a private non-negative value $v_i \in \mathbb{R}_+$ for winning, and the bid they submit consists of a single number $b_i \in \mathbb{R}_+$. As before, we assume *quasi-linear* utilities: the utility of winning is $v_i - \text{payment}_i$ and the utility of losing is $-\text{payment}_i$.

Examples of single-parameter auction are when we have one item to sell, or k copies to sell and each player can win at most one copy. However, the most general form is the following:

- For each vector b of n bids from the players, the mechanism declares some winners and the rest losers, and charges some payments to each player.

So formally, a mechanism is a function $\mathbb{R}^n_+ \to \mathbf{2}^n \times \mathbb{R}^n_+$ where the inputs are the bids and the outputs are the winners and payments. We'll write

$$\text{winners}(b) \subseteq \{1, 2, \ldots, n\}$$

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

to denote the players who win for a given set of bids, in a given mechanism. As usual we are interested in truthful auctions. It turns out for truthful mechanisms the winners(b) function will also implicitly determine the prices!

The first result we will talk about can be thought of as a warm-up. A mechanism is said to be *normalized* if losers always pay 0; this assumption is intuitive, and also without loss of generality if we care only about Nash equilibria, weak dominance, or truthfulness (using the h_{-i} lemma the from last lecture). We call a mechanism *voluntary* if every player who bids 0 pays 0.

Theorem 1. *A voluntary normalized mechanism (for single-parameter bidders) is truthful if and only if it has the following properties.*

monotony *Whenever bid vector b causes player i to win, so does any bid vector (b'_i, b_{-i}) such that $b'_i > b_i$.*

payment rule *If player i wins an item when bid vector b is submitted, then* payment$_i$ *equals the* critical bid

$$\inf\{b'_i \geq 0 \mid i \in \text{winners}(b'_i, b_{-i})\}.$$

Proof. First, let us see why monotony is necessary for the mechanism to be truthful (i.e., for $b_i = v_i$ to be weakly dominant). Suppose the mechanism is truthful but for the sake of contradiction that monotony is violated for some $b_{-i}, b_i < b'_i : i \in \text{winners}(b)$ and $i \notin \text{winners}(b'_i, b_{-i})$. Let p be the payment made by player i when they win in b. Fix the bids of the opponents at b_{-i}. For a player i with true value $v_i = b_i$, truthfulness implies

$$\underbrace{b_i - p}_{\text{utility with truthful bid}} \geq \underbrace{0.}_{\text{utility with bid } b'_i}$$

Similarly for a player with true value b'_i, we get $0 \geq b'_i - p$. This gives $b_i \geq b'_i$, a contradiction.

Next, why is the payment rule necessary? Note that for any fixed b_{-i}, monotony implies the set of bids b_i for which $i \in \text{winners}(b)$ is a subset of $\mathbb{R}_+$ which is "up-closed:" it is $\varnothing$ or an infinite interval of the form $(x, +\infty)$ or $[x, +\infty)$. The payment for player i under two different winning bids cannot change, since a truthful bidder paying the higher price would have incentive to lie and report the smaller one. Finally, the constant price p paid in all winning situations can neither be higher than x (or else a bidder with $x < v_i < p$ would have incentive to lie and report 0, using voluntariness) nor lower than x (or similarly a bidder with $p < v_i < x$ would have incentive to lie and report a winning bid).

Finally, we prove the converse: every monotonic voluntary normalized mechanism with critical bid payments is truthful. Every winner gets a non-negative utility (since by the payment rule their payment is no more than their bid=value), and deviating can only cause them to lose, getting utility 0; every loser gets utility 0 and deviating can only cause them to win, similarly getting nonpositive utility. □

Voluntariness can be removed if one is willing to make the theorem a tiny bit more complicated; see Theorem 9.36 of Nisan et al. A new possibility arises: a player could always win an item and pay some fixed positive price.

Assumption 1. *In the rest of the lecture, we will assume all mechanisms are normalized and voluntary.*

The *income* in an auction is the sum of the payments.

Exercise. Consider auctioning a single item (i.e. there is always *exactly* one winner) to two bidders. The VCG auction/second-price auction awards the item to the highest bidder (breaking ties arbitrarily), and charges them the second-highest bid $b^{(2)}$. Show that for any truthful voluntary normalized mechanism, if its income is at least $b^{(2)}$ for all b, then the mechanism is a second-price auction.

Exercise. Now consider auctioning two items (we want $|\text{winners}(b)| = 2$ for all b), this time to three players. The VCG mechanism awards the items to the two highest bidders, and charges them each the third-highest price $b^{(3)}$. Find a truthful voluntary normalized mechanism such that (i) for all b, its income is at least $2b^{(3)}$ (ii) for some b, the winners don't correspond to the top two bids (so unlike VCG it does not maximize social welfare).

1.1 Maximizing Income

With Theorem 1 we can start to address the type of problem which is the main subject today.

Example 2. As a warm-up, we consider the case of just a single bidder and selling just a single item. This bidder arrives with a valuation/bid b drawn from some distribution F (we conflate bids with true values since we will only look at truthful mechanisms.) We want to design a single-parameter mechanism (namely, either we sell them the item, or we don't) which maximizes our expected income (with respect to the bidder's probability distribution).

Theorem 1 implies that our mechanism is truthful so long as the set of winning bids is up-closed (by monotony). If we define the set of winning bids as $[x, +\infty)$ then we have to sell the item at price x, by the payment rule. Therefore the expected income is

$$x \cdot \Pr[b \geq x]$$

and we simply pick x which maximizes this quantity. (Note setting the winning bids to a set of the form $(x, +\infty)$ is also possible but never better.) We call x the *optimal monopoly price.*

The resulting auction is the same as a one-player second-price auction with reserve price x. It is useful to look at one concrete example: suppose the distribution F is uniform on $[0, 1]$, then the expected income for reserve price x is

$$x \cdot \Pr[b \geq x] = x(1 - x)$$

which is maximized at $x = 1/2$, with expected income $1/4$.

To motivate the next part, now consider two bidders arriving, both with independent random bids drawn uniformly from $[0, 1]$, where we want a truthful mechanism to sell just a single item. Write b_1, b_2 for the bids. If we use a second-price auction, we get income $\min\{b_1, b_2\}$ and thus the expected income is

$$\int_0^1 \int_0^1 \min\{b_1, b_2\} db_1 db_2 = 1/3.$$

On the other hand, adding the same reserve price of 1/2 (from the single-bidder case) actually increases the expected income to

$$\int_0^{1/2} \int_0^{1/2} 0 db_1 db_2 + \int_0^{1/2} \int_{1/2}^1 \frac{1}{2} db_1 db_2 + \int_{1/2}^1 \int_0^{1/2} \frac{1}{2} db_1 db_2 + \int_{1/2}^1 \int_{1/2}^1 \min\{b_1, b_2\} db_1 db_2 = \frac{5}{12}.$$

What happened is that we lost some money some of the time (when both players have bids less than 1/2) but gained even more money other times. This $\frac{5}{12}$ turns out to be the *best* we can do for two-player auctions — even if rather than a truthful deterministic mechanism we allow a randomized mechanism that is truthful-in-expectation. We will aim to prove this in a general setting.

2 Truthful-in-Expectation Mechanisms

Myerson's results consider truthful-in-expectation auctions, when players have valuations coming from independent random distributions. Here is one useful tool.

Lemma 3. *Fix a (voluntary normalized) mechanism. Fix a distribution on bid vectors for player i's opponents. As a function of player i's bid b_i, let $0 \leq x_i(b_i) \leq 1$ be the probability that player i is a winner. The mechanism is truthful in expectation to player i (w.r.t. the distribution on b_{-i}) if and only if*

monotony $x_i(b_i)$ *is a weakly monotonically increasing function*

payment rule $\mathsf{E}[\text{payment}_i(b_i)] = \int_0^{b_i} (x_i(b_i) - x_i(z)) dz.$

The salient points are that the interaction with other players only shows up in $x_i(b_i)$, and that there is a formula for payments. We won't give the proof since it's very similar to Theorem 1; it also reduces to Theorem 1 in the case that b_{-i} is a singleton distribution and x is integer-valued. The complete proof is given as Theorem 9.39 in Nisan et al.; it also holds if we allow the mechanism to have some randomness of its own (so you can actually take the randomness in b_{-i} and move it into the mechanism).

From now on, we think of each player i's bids/true values as being drawn from a distribution F_i. Our goal is to determine, among all truthful voluntary normalized mechanisms, the one which maximizes the expected income. Myerson's results won't exactly give us the whole picture, but they work in a pretty broad setting which we will make precise as we go along. There is an *ironing* procedure to try to fix distributions not meeting our assumptions but we don't discuss it in detail. We will have two sets of assumptions: here is the first one.

Definition 4. We write $F_i(x) := \Pr_{b_i \sim F_i}[b_i \leq x]$ for the cumulative density function of F_i. The probability density function, denoted $f_i(x)$, is $F_i'(x)$.

Assumption 2. *F_i is continuous, differentiable, and the range of values with positive probability density is an interval $(0, h)$.*[1]

Definition 5. The *virtual valuation* of player i with value b_i is

$$\phi_i(b_i) := b_i - \frac{1 - F_i(b_i)}{f_i(b_i)}.$$

These definitions are valuable because of the following, which we will prove later, and which extends Lemma 3:

Lemma 6. *Fix a distribution for b_{-i} and so define $x(b_i)$, as before. Fix a distribution F_i for player i, which defines ϕ_i. Assuming truthfulness, the expected payment of player i equals the expected value of $\phi_i(b_i)x_i(b_i)$.*

Given our definition of *virtual valuation*, and noting that $\sum_i b_i x_i$ is the expected social welfare, it makes sense to call

$$\mathsf{E}[\sum_i \phi_i(b_i)x_i(b_i)]$$

the *expected virtual social welfare.* Then adding Lemma 6 over all players gives:

Corollary 7. *The expected profit of a mechanism equals its expected virtual social welfare (for independently distributed bidders, under our assumptions).*

Moreover, we have a truthful mechanism which maximizes social welfare (VCG)! Once we have this fact, it is tempting to try running VCG on the virtual valuations:

Virtual-VCG

1. Given the bids b, compute the virtual bids $\phi_i(b_i)$ for each i
2. Run VCG on the virtual bids:

   pick the feasible winner set W maximizing $\sum_{i \in W} \phi_i(b_i)$

   compute virtual prices p_i' for winners, using the VCG formula on the virtual bids
3. Charge each winner i the price $\phi_i^{-1}(p_i')$

[1] The continuity is not essential but makes our calculations more straightforward; we could relax differentiability to differentiability "almost everywhere;" the last part "no holes" is essential and without it Lemma 6 would be false, in short since we could get two different x_i that behave the same on the range of b_i.

Before proceeding further, it is helpful to return to the example $F_i = \text{uniform}(0, 1)$ for all i when we want to auction a single item. The probability density function is a constant, 1. So the virtual valuation of b_i is $\phi_i(b_i) = b_i - (1 - b_i)/1 = 2b_i - 1$. Running VCG — maximizing the virtual social welfare — for an auction on the virtual bids would award the items to the top virtual bid (or to none, if all are negative). The winner would be charged the 2nd largest bid (or $\phi^{-1}(0) = 1/2$ if only one is positive). This is equivalent to the VCG-with-reserve-1/2 studied earlier! Thus we have proven among all truthful-in-expectation mechanisms for these bidders, the one maximizing the income is this simple deterministic one.

Exercise. What happens if we have k identical items to auction away (and each player can win only one)?

However, this special case is not totally representative. It can be that VIRTUAL-VCG is not truthful! We have shown that VCG is truthful on its inputs, but here the inputs are the *virtual* bids.

Assumption 3. *The function $\phi_i(b_i)$ is monotonically increasing in b_i.*

This assumption is a special case of a so-called *monotone hazard rate* from economics; see Section 13.2.2 of Nisan et al. If we assume this, monotony of x in virtual bids is equivalent to monotony of x in real bids. Consequently, using Theorem 1 and the fact that VCG satisfies monotony,

Theorem 8. *Under our assumptions,* VIRTUAL-VCG *gives the maximum possible expected income among all truthful-in-expectation mechanisms.*

To prove this, it only remains to prove Lemma 6.

Proof of Lemma 6. We want to show $\mathsf{E}[\text{payment}_i] = \mathsf{E}[\phi_i x_i]$. Shorten $\text{payment}_i(b_i)$ to $p_i(b_i)$, and let us drop all i subscripts for convenience.

When b is drawn according to F, calculus tells us the expected value of any function $g(b)$ equals the integral $\int_{b=0}^{h} g(b)f(b)db$. We likewise have that $1-F(u) = \Pr[b > u] = \int_{z=u}^{h} f(z)dz$.

Thus we want to show that

$$\int_{b=0}^{h} p(b)f(b)db = \int_{b=0}^{h} x(b)\phi(b)f(b)db. \tag{1}$$

We expand ϕ using its definition, and p using the formula from Lemma 3.

$$\begin{aligned}
(1) \Leftrightarrow \int_{b=0}^{h}\int_{z=0}^{b} l(x(b) - x(z))dzf(b)db &= \int_{b=0}^{h} x(b)\left(b - \frac{1-F(b)}{f(b)}\right)f(b)db.\\
\Leftrightarrow \int_{b=0}^{h} bx(b)f(b)db - \int_{b=0}^{h}\int_{z=0}^{b} x(z)f(b)db &= \int_{b=0}^{h} bx(b)f(b)db - \int_{b=0}^{h}(1-F(b))x(b)db\\
\Leftrightarrow \int_{b=0}^{h}\int_{z=0}^{b} x(z)f(b)db &= \int_{b=0}^{h}\int_{z=b}^{h} f(z)x(b)db
\end{aligned}$$

and the last line is true by exchanging the order of integration. □

Exercise. Suppose the auctioneer attaches a value $v_0 > 0$ to keeping the item: so instead of revenue they want to maximize their (expected) *surplus*, which equals income if they sell them item, and equals v_0 otherwise. Under the same assumptions as above, what truthful mechanism maximizes the expected surplus?

A remark: we assumed that the auctioneer has some valid winner sets, and others which are invalid. Myerson's results actually apply to the case that every winner set $W \subseteq \{1, \ldots, n\}$ is possible but each has a cost $c(W)$ to the auctioneer. (The valid-invalid setting is equivalent to $c(W) \in \{0, +\infty\}$ for all W). What we find is that (under our assumptions) the truthful-in-expectation mechanism which maximizes $\mathsf{E}[\text{income - } c(W)]$ is a modified virtual VCG, where the virtual social welfare of W is $\sum_{i \in W} \phi_i(b_i) - c(W)$ (and the payment rule is modified similarly). For example, if $c(W) = |W|$, the optimal truthful mechanism makes every player with $b_i \geq 1$ a winner, and charges them 1.

Exercise. Myerson's result proves that for the class of F's meeting all our assumptions, the revenue-maximizing auction for a single item among n players is a second-price auction with reserve x, where x does not depend on n. Show that this conclusion does not hold for *all* distributions. (Hint: one possible counterexample is a discrete distribution on a small number of values, where the optimal reserve price for 1 player is not equal to the optimal reserve price for 2 players.)

Exercise. (Bulow-Klemperer [1]) We will compare two auctions. You can perform the first three parts independently of one another.

1. Under both assumptions, show that the expected virtual valuation for player i is 0.
2. Show for any two independent random variables x, y, $\mathsf{E}[\max\{x, y\}] \geq \mathsf{E}[\max\{x, \mathsf{E}[y]\}]$.
3. Take F as described above; assume all bidders have independent valuations drawn according to F. Consider two second-price auctions: (i) we auction one item to n bidders, using the optimal reserve price; and (ii) we auction one item to $n + 1$ bidders, using no reserve price. Show that the expected income of (ii) is at least as large as that of (i). (Moral: rather than be smart and compute the optimal reserve price, it's better to just attract one more bidder and use no reserve price.)
4. Generalize the previous result to k-item auctions: how many more bidders must you attract?

References

[1] J. Bulow and P. Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):pp. 180–194, 1996.

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