Lecture 6

Optimal Auctions Continued

1 Recap

Last week, we...

- Set up the Myerson auction environment:
  - n risk-neutral bidders
  - independent types $t_i \sim F_i$ with support $[a_i, b_i]$
  - residual valuation of t_0 for the seller
- Named our goal: maximize expected seller payoff over all conceivable auction/sales formats, subject to two constraints: bidders participate voluntarily and play equilibrium strategies
- Defined mechanisms, and showed the Revelation Principle – that without loss of generality, we can restrict attention to direct revelation mechanisms
- Showed that feasibility of a mechanism is equivalent to four conditions holding:
  4. Q_i is weakly increasing in t_i
  5. $U_i(t_i) = U_i(a_i) + \int_{a_i}^{t_i} Q_i(p, s) ds$ for all i, all t_i
  6. $U_i(a_i) \geq 0$ for all i, and
  7. $\sum_j p_j(t) \leq 1$, $p_i \geq 0$
- So we defined our goal as solving

$$\max_{direct\ \ revelation\ \ mechanisms} E_t \left\{ \sum_{i \in N} x_i(t) + t_0 \left(1 - \sum_{i \in N} p_i(t) \right) \right\} \qquad s.t. \qquad (4), (5), (6), (7)$$

- So... onward!

2 Rewriting our Objective Function

By adding and subtracting $E_i \sum_i p_i(t)t_i$, we can rewrite the seller's objective function as

$$\begin{aligned} U_0(p,x) &= E_t \left\{ \sum_{i \in N} x_i(t) + t_0 \left(1 - \sum_{i \in N} p_i(t) \right) \right\} \\ &= t_0 + \sum_{i \in N} E_t p_i(t)(t_i - t_0) - \sum_{i \in N} E_t(p_i(t)t_i - x_i(t)) \end{aligned}$$

Note that...

- the first term is a constant (the payoff to doing nothing)
- the second term is the total surplus created by selling the object
- the third term, which is being subtracted, is the expected payoff going to the bidders

Next, we do some work expressing the last term, expected bidder surplus, in a different form:

$$\begin{aligned} E_t(p_i(t)t_i - x_i(t)) &= \textstyle\int_{a_i}^{b_i} U_i(t_i) f_i(t_i) dt_i \\ &= \textstyle\int_{a_i}^{b_i} \left(U_i(a_i) + \int_{a_i}^{t_i} Q_i(s_i) ds_i \right) f_i(t_i) dt_i \\ &= U_i(a_i) + \textstyle\int_{a_i}^{b_i} \int_{a_i}^{t_i} Q_i(s_i) f_i(t_i) ds_i dt_i \\ &= U_i(a_i) + \textstyle\int_{a_i}^{b_i} \int_{s_i}^{b_i} f_i(t_i) Q_i(s_i) dt_i ds_i \\ &= U_i(a_i) + \textstyle\int_{a_i}^{b_i} (1 - F_i(s_i)) Q_i(s_i) ds_i \\ &= U_i(a_i) + \textstyle\int_{a_i}^{b_i} (1 - F_i(s_i)) \left(\int_{t_{-i}} p_i(s_i, t_{-i}) f_{-i}(t_{-i}) dt_{-i} \right) ds_i \\ &= U_i(a_i) + \textstyle\int_{a_i}^{b_i} \frac{1-F_i(s_i)}{f_i(s_i)} \left(\int_{t_{-i}} p_i(s_i, t_{-i}) f_{-i}(t_{-i}) dt_{-i} \right) f_i(s_i) ds_i \\ &= U_i(a_i) + \textstyle\int_T \frac{1-F_i(s_i)}{f_i(s_i)} p_i(t_{-i}, s_i) f(t_{-i}, s_i) dt_{-i} ds_i \\ &= U_i(a_i) + \textstyle\int_T \frac{1-F_i(t_i)}{f_i(t_i)} p_i(t) f(t) dt \\ &= U_i(a_i) + E_t p_i(t) \frac{1-F_i(t_i)}{f_i(t_i)} \end{aligned}$$

- So we can rewrite the objective function as

$$\begin{aligned} U_0(p,x) &= t_0 + \sum_{i \in N} E_t p_i(t)(t_i - t_0) - \sum_{i \in N} U_i(a_i) - \sum_{i \in N} E_t p_i(t) \frac{1 - F_i(t_i)}{f_i(t_i)} \\ &= t_0 + \sum_{i \in N} E_t p_i(t) \left(t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} - t_0 \right) - \sum_{i \in N} U_i(a_i) \end{aligned}$$

- So the seller's problem amounts to choosing an allocation rule p and expected payoffs for the low types $U_i(a_i)$ to maximize

$$t_0 + E_t \left\{ \sum_{i \in N} p_i(t) \left(t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} - t_0 \right) \right\} - \sum_{i \in N} U_i(a_i)$$

  subject to feasibility – which just means p plausible, Q_i increasing in type, and $U_i(a_i) \geq 0$

- (The envelope formula for bidder payoffs is no longer a constraint – we've already imposed it.)

- Once we find a mechanism we like, each U_i is uniquely determined by the envelope formula, and so the rest of the transfers x_i are set to satisfy those required payoffs.

- Once we've phrased the problem in this way, Myerson points out, **revenue equivalence** becomes a straightforward corollary:

Corollary 1. *For a given environment, the expected revenue of an auction depends only on the equilibrium allocation rule and the expected payoffs of the lowest possible type of each bidder.*

The Revenue Equivalence Theorem is usually stated in this way:

Corollary 2. *Suppose bidders have symmetric independent private values and are risk-neutral. Define a* standard auction *as an auction where the following two properties hold:*

1. *In equilibrium, the bidder with the highest valuation always wins the object*

2. *The expected payment from a bidder with the lowest possible type is* 0

Any two standard auctions give the same expected revenue.

Two standard auctions also give the same expected surplus to each type of each bidder $U_i(t_i)$.

So this means with symmetry, independence, and risk-neutrality, *any* auction with a symmetric, strictly-monotone equilibrium gives the same expected revenue. (Examples.)

Now, back to maximizing expected revenue

- We've redefined the problem as choosing p and $U_i(a_i)$ to maximize

$$t_0 + E_t \left\{ \sum_{i \in N} p_i(t) \left(t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} - t_0 \right) \right\} - \sum_{i \in N} U_i(a_i)$$

- Clearly, to maximize this, we should set $U_i(a_i) = 0$ for each i
- This leaves Myerson's Lemma 3:

Lemma 1. *If p maximizes*

$$t_0 + E_t \left\{ \sum_{i \in N} p_i(t) \left(t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} - t_0 \right) \right\}$$

subject to Q_i increasing in t_i and p possible, and

$$x_i(t) = t_i p_i(t) - \int_{a_i}^{t_i} p_i(t_{-i}, s_i) ds_i$$

then (p, x) is an optimal auction.

- The transfers are set to make $U_i(a_i) = 0$ and give payoffs required by the envelope theorem
- To see this, fix t_i and take the expectation over t_{-i}, and we find

$$E_{t_{-i}} x_i(t) = t_i Q(t_i) - \int_{a_i}^{t_i} Q_i(s_i) ds_i$$

  or

$$\int_{a_i}^{t_i} Q_i(s_i) ds_i = t_i Q_i(t_i) - E_{t_{-i}} x_i(t_i, t_{-i}) = U_i(t_i)$$

  which is exactly the envelope theorem combined with $U_i(a_i) = 0$

- (The exact transfers $x_i(t)$ are not uniquely determined by incentive compatibility and the allocation rule p; what is uniquely pinned down is $E_{t_{-i}} x_i(t_i, t_{-i})$, because this is what's payoff-relevant to bidder i. The transfers above are just one rule that works.)
- Next, we consider what the solution looks like for various cases.

3 The Regular Case

- With one additional assumption, things fall into place very nicely.
- Define a distribution F_i to be *regular* if

$$t_i - \frac{1 - F_i(t_i)}{f_i(t_i)}$$

is strictly increasing in t_i

  - This is not that crazy an assumption
  - Most familiar distributions have increasing hazard rates – that is, $\frac{f}{1-F}$ is increasing, which would imply $\frac{1-F}{f}$ decreasing
  - This is a weaker condition, since $\frac{f}{1-F}$ is allowed to decrease, just not too quickly.

- When the bid distributions are all regular, the optimal auction becomes this:
  - Calculate $c_i(t_i) = t_i - \frac{1-F_i(t_i)}{f_i(t_i)}$ for each player
  - If $\max_i c_i(t_i) < t_0$, keep the good; if $\max_i c_i(t_i) \geq t_0$, award the good to the bidder with the highest value of $c_i(t_i)$
  - Charge the transfers determined by incentive compatibility and this allocation rule

- Note that this rule makes Q_i monotonic – $Q_i(t_i) = 0$ for $t_i < c_i^{-1}(t_0)$, and $\prod_{j \neq i} F_j \left(c_j^{-1}(c_i(t_i)) \right)$ above it
- So the rule satisfies our two constraints, and it's obvious that it maximizes the seller's objective function
- There's even a nice interpretation of the x_i we defined above. Fixing everyone else's type, p_i is 0 when $c_i(t_i) < \max\{t_0, \max\{c_j(t_j)\}\}$ and 1 when $c_i(t_i) > \max\{t_0, \max\{c_j(t_j)\}\}$, so this is just

$$x_i(t) = t_i - \int_{t_i^*}^{t_i} ds_i = t_i - (t_i - t_i^*) = t_i^*$$

where t_i^* is the lowest type that i could have reported (given everyone elses' reports) and still won the object.

- This payment rule makes incentive-compatibility obvious: for each combination of my opponents' bids, I face some cutoff t_i^* such that if I report $t_i > t_i^*$, I win and pay t_i^*; and if I report less than that, I lose and pay nothing. Since I want to win whenever $t_i > t_i^*$, just like in a second-price auction, my best-response is to bid my type.

3.1 Symmetric IPV

In the case of symmetric IPV, each bidder's c function is the same as a function of his type, that is,

$$c_i(t_i) = c(t_i) = t_i - \frac{1 - F(t_i)}{f(t_i)}$$

which is strictly increasing in t_i. This means the bidder with the highest type has the highest c_i, and therefore gets the object; and so his payment is the reported type of the next-highest bidder, since this is the lowest type at which he would have won the object. Which brings us to our first claim:

Theorem 1. *With symmetric independent private values, the optimal auction is a second-price auction with a reserve price of* $c^{-1}(t_0)$*.*

Note, though, that even when $t_0 = 0$, this reserve price will be positive. The optimal auction is not efficient – since $c(t_i) < t_i$, the seller will sometimes keep the object even though the highest bidder values it more than him – but he never allocates it to the "wrong" bidder.

Also interesting is that the optimal reserve price under symmetric IPV does not depend on the number of bidders – it's just $c^{-1}(t_0)$, regardless of N.

3.2 Asymmetric IPV

When the bidders are not symmetric, things are different. With different F_i, it will not always be true that the bidder with the highest c_i also has the highest type; so sometimes the winning bidder will not be the bidder with the highest value. (As we'll see later, efficiency is not standard in auctions with asymmetric bidders: even a standard first-price auction is sometimes not be won by the bidder with the highest value.)

One special case that's easy to analyze: suppose every bidder's bid is drawn from a uniform distribution, but uniform over different intervals. That is, suppose each F_i is the uniform distribution over a (potentially) different interval $[a_i, b_i]$. Then

$$c_i(t_i) = t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} = t_i - \frac{(b_i - t_i)/(b_i - a_i)}{1/(b_i - a_i)} = t_i - (b_i - t_i) = 2t_i - b_i$$

So the optimal auction, in a sense, penalizes bidders who have high maximum valuations. This is to force them to bid more aggressively when they have high values, in order to extract more revenue; but the price of this is that sometimes the object goes to the wrong bidder.

3.3 What About The Not-Regular Case?

Myerson does solve for the optimal auction in the case where c_i is not increasing in t_i, that is, where the auction above would not be feasible. Read the paper if you're interested.

4 Bulow and Klemperer, "Auctions versus Negotiations"

We just learned that with symmetric IPV and risk-neutral bidders, the best you can possibly do is to choose the perfect reserve price and run a second-price auction. This might suggest that choosing the perfect reserve price is important. There's a paper by Bulow and Klemperer, "Auctions versus Negotiations," that basically says: nah, it's not that important. Actually, what they say is, it's better to attract one more bidder than to run the perfect auction.

Suppose we're in a symmetric IPV world where bidders' values are drawn from some distribution F on $[a, b]$, and the seller values the object at t_0. Bulow and Klemperer show the following: as long as $a \geq t_0$ (all bidders are "serious"), the optimal auction with N bidders gives lower revenue than a second-price auction with no reserve price and $N+1$ bidders.

To see this, recall that we wrote the auctioneer's expected revenue as

$$t_0 + E_t \left\{ \sum_{i \in N} p_i(t) \left(t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} - t_0 \right) \right\} - \sum_{i \in N} U_i(a_i)$$

Consider mechanisms where $U_i(a_i) = 0$, and define the **marginal revenue** of bidder i as

$$MR_i = t_i - \frac{1 - F_i(t_i)}{f_i(t_i)}$$

so expected revenue is

$$E_t \left\{ \sum_{i \in N} p_i(t) MR_i(t) + \left(1 - \sum_{i \in N} p_i(t) \right) t_0 \right\}$$

So if we think of the seller as being another possible buyer, with marginal revenue of t_0, then the expected revenue is simply the expected value of the marginal revenue of the winner.

Jump back to the symmetric case, so $F_i = F$. Continue to assume regularity. In an ordinary second-price or ascending auction, with no reserve price, the object sells to the bidder with the highest type, which is also the bidder with the highest marginal revenue; so the expected revenue in this type of auction (what Bulow and Klemperer call an "absolute English auction") is

$$Expected \;\; Revenue = E_t \max\{MR(t_1), MR(t_2), \ldots, MR(t_N)\}$$

(This is Bulow and Klemperer Lemma 1.)

The fact that expected revenue = expected marginal revenue of winner also makes it clear why the optimal reserve price is $MR^{-1}(t_0)$ – this replaces bidders with marginal revenue less than t_0 with t_0. So (counting the seller's value from keeping the unsold object) an English auction with an optimal reserve price has expected revenue

$$Expected \;\; Revenue = E_t \max\{MR(t_1), MR(t_2), \ldots, MR(t_N), t_0\}$$

So here's the gist of Bulow and Klemperer, "Auctions Versus Negotiations." They compare the simple ascending auction with $N+1$ bidders, to the optimal auction with N bidders. (We discovered last week that with symmetric independent private values, the optimal auction is an ascending auction with a reserve price of $MR^{-1}(t_0)$.) The gist of Bulow and Klemperer is that the former is higher, that is, that

$$E\max\{MR(t_1), MR(t_2), \dots, MR(t_N), MR(t_{N+1})\} \geq E\max\{MR(t_1), MR(t_2), \dots, MR(t_N), t_0\}$$

so the seller gains more by attracting one more bidder than by holding the "perfect" auction. (They normalize t_0 to 0, but this doesn't change anything.)

Let's prove this. The proof has a few steps.

First of all, note that the expected value of $MR(t_i)$ is a, the lower bound of the support. This is because

$$\begin{aligned} E_{t_i} MR(t_i) &= E_{t_i}\left(t_i - \tfrac{1-F(t_i)}{f(t_i)}\right) \\ &= \int_a^b \left(t_i - \tfrac{1-F(t_i)}{f(t_i)}\right) f(t_i) dt_i \\ &= \int_a^b \left(t_i f(t_i) - 1 + F(t_i)\right) dt_i \end{aligned}$$

Now, $tf(t) + F(t)$ has integral $tF(t)$, so this integrates to

$$t_i F(t_i)|_{t_i=a}^{b} - (b-a) = b - 0 - (b-a) = a$$

which by assumption is at least t_0. So $E(MR(t_i)) \geq t_0$.

Next, note that for fixed x, the function $g(y) = \max\{x, y\}$ is convex, so by Jensen's inequality,

$$E_y \max\{x, y\} \geq \max\{x, E(y)\}$$

If we take an expectation over x, this gives us

$$E_x \{E_y \max\{x, y\}\} \geq E_x \max\{x, E(y)\}$$

or

$$E\max\{x, y\} \geq E\max\{x, E(y)\}$$

Now let $x = \max\{MR(t_1), MR(t_2), \dots, MR(t_N)\}$ and $y = MR(t_{N+1})$;

$$E\max\{MR(t_1), MR(t_2), \dots, MR(t_N), MR(t_{N+1})\} \geq$$

$$E\max\{MR(t_1), MR(t_2), \dots, MR(t_N), E(MR(t_{N+1}))\} \geq$$

$$E\max\{MR(t_1), MR(t_2), \dots, MR(t_N), t_0\}$$

and that's the proof.

Finally (and leading to the title of the paper), Bulow and Klemperer point out that "negotiations" – really, any process for allocating the object and determining the price – cannot outperform the optimal mechanism, and therefore leads to lower expected revenue than a simple ascending auction with one more bidder. They therefore argue that a seller should never agree to an early "take-it-or-leave-it" offer from one buyer when the alternative is an ascending auction with at least one more buyer, etc.

References

- Roger Myerson (1981), "Optimal Auctions," *Mathematics of Operations Research* 6
- Jeremy Bulow and Paul Klemperer (1996), "Auctions Versus Negotiations," *American Economic Review* 86