CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

Allan Borodin (instructor)
Tyrone Strangway and Young Wu (TAs)

October 19, 2016

Lecture 11

- Announcements
  - I have posted the first two questions for Assignment 2 which is due October 28. I expect to post another question later tonday or tonight.
  - Assignment 1 has been graded. Some statistics:
    average 73.9%, median 78.9%, 6 below 50%
    One person did not submit via Markus.
- Todays agenda
  We continue the discussion of auctions in Chapter 14.
  - We continue to give a very fast overview of some of the important results in Chapter 14 of the KP text. Again, I indicate that this is not an easy chapter to read and we will try to come back to these results as we proceed with our discussion of auctions in this and other chapters.
  - Where we left off last time ... the revenue equivalence theorem
  - Individual rationality
  - The revelation principle
  - The main result to be discussed today is Myerson's optimal algorithm

Revenue Equivalence Theorem

Here again is a statement given in Vijay Krishna's "Auction Theory" text. (See also, the informal statement in Wikipedia and a formal statement in Theorem 14.4.2).

Revenue Equivalence Theorem

Here again is a statement given in Vijay Krishna's "Auction Theory" text. (See also, the informal statement in Wikipedia and a formal statement in Theorem 14.4.2).

The Revenue Equivalence Theorem for selling a single item

Suppose that values are independently and identically distributed and all bidders are risk neutral. Then any symmetric and increasing equilibrium of any standard auction, such that the expected payment of a bidder with value zero is zero, yields the same expected revenue to the seller.

Revenue Equivalence Theorem

Here again is a statement given in Vijay Krishna's "Auction Theory" text. (See also, the informal statement in Wikipedia and a formal statement in Theorem 14.4.2).

The Revenue Equivalence Theorem for selling a single item

Suppose that values are independently and identically distributed and all bidders are risk neutral. Then any symmetric and increasing equilibrium of any standard auction, such that the expected payment of a bidder with value zero is zero, yields the same expected revenue to the seller.

The proof given in the text assumes that $a(w) = F^{n-1}$ is differentiable but doesn't say continuous. However it does say that that the cumulative distribution function F is strictly increasing which would then not apply to discrete probability distributions having a step function. But ....

Can the revenue equivalence theorem be extended?

For a one item auction when all buyers are drawing values from an i.i.d distribution, the revenue equivalence theorem is quite general.

As the KP text says, it can be extended to deal with randomized pricing and reserve pricing.

And as the text also says, it can be extended to the case of k identical items (priced the same) and "unit demand" buyers.

Can the revenue equivalence theorem be extended?

For a one item auction when all buyers are drawing values from an i.i.d distribution, the revenue equivalence theorem is quite general.

As the KP text says, it can be extended to deal with randomized pricing and reserve pricing.

And as the text also says, it can be extended to the case of k identical items (priced the same) and "unit demand" buyers.

It is natural to ask if the theorem can be extended beyond the continuous i.i.d case or for more than one item.

- Having asked someone more knowledgeable than me, I am told that the theorem can be extended to i.i.d. discrete distributions although in the Economics literature, it seems that one is always considering continuous distributions. I am told that proving the theorem for discrete i.i.d. distributions would "require care".
- The theorem *does not* extend in general to the asymmetric case! Nor does it extend to the case of muitiple items.

An example of the asymmetric case

The following is an example (Krishna, Example 4.3) shows how the expected revenue of a 1st price auction can exceed that of a 2nd price auction.

Let the value of buyer 2 be drawn from $U[0, \frac{1}{1+\alpha}]$ while the value of buyer 1 is drawn from $U[0, \frac{1}{1-\alpha}]$ some $0 < \alpha < 1$.

Krishna shows that the equilibrium prices in a 2nd price auction is less than if $\alpha = 0$ (i.e. the distribution $U[0, 1]$) whereas for a 1st price auction the equibrium prices are higher than for the $U[0, 1]$ distribution.

As the KP text explains, buyer 1 (with a smaller possible maximum valuation) has to bid more agressively than bidder 2. This suggests that there will be valuations $v_1 < v_2$ such that $\beta_1(v_1) > \beta_2(v_2)$ so that buyer 1 wins the auction. But note that β_1 and β_2 are both differentiable and strictly increasing.

Bayes-Nash equlibrium and the price of anarchy

While revenue equivalence need not hold in asymmetric distributions, there is a Bayes- Nash equilibrium and it is characterized in Theorem 14.6.1.

Bayes-Nash equlibrium and the price of anarchy

While revenue equivalence need not hold in asymmetric distributions, there is a Bayes- Nash equilibrium and it is characterized in Theorem 14.6.1.

Aside: It is easy to misread Theorem 14.6.1 to think that it is suggesting revenue equivalence.

Bayes-Nash equlibrium and the price of anarchy

While revenue equivalence need not hold in asymmetric distributions, there is a Bayes- Nash equilibrium and it is characterized in Theorem 14.6.1.

Aside: It is easy to misread Theorem 14.6.1 to think that it is suggesting revenue equivalence.

Since a 2nd price auction insures truthfulness, we can assume that buyers bid their values and hence the auction is won by the buyer with the highest value. This insures that the 2nd price auction is optimal with respect to social welfare (i.e. price of anarchy $POA = 1$).

Bayes-Nash equlibrium and the price of anarchy

While revenue equivalence need not hold in asymmetric distributions, there is a Bayes- Nash equilibrium and it is characterized in Theorem 14.6.1.

Aside: It is easy to misread Theorem 14.6.1 to think that it is suggesting revenue equivalence.

Since a 2nd price auction insures truthfulness, we can assume that buyers bid their values and hence the auction is won by the buyer with the highest value. This insures that the 2nd price auction is optimal with respect to social welfare (i.e. price of anarchy $POA = 1$).

In a 1st price auction, the buyer with the highest value may not win the auction. So is there any bound on the POA for a 1st price auction?

Bayes-Nash equlibrium and the price of anarchy

While revenue equivalence need not hold in asymmetric distributions, there is a Bayes- Nash equilibrium and it is characterized in Theorem 14.6.1.

Aside: It is easy to misread Theorem 14.6.1 to think that it is suggesting revenue equivalence.

Since a 2nd price auction insures truthfulness, we can assume that buyers bid their values and hence the auction is won by the buyer with the highest value. This insures that the 2nd price auction is optimal with respect to social welfare (i.e. price of anarchy $POA = 1$).

In a 1st price auction, the buyer with the highest value may not win the auction. So is there any bound on the POA for a 1st price auction?

Theorem: In a Bayes-Nash equilibrium for selling one item by a 1st price auction, $POA \geq 1/2$; that is, the winning bidder obtains at least $1/2$ of the maximum value amongst all the buyers. The buyers distributions can be non identical and even correlated

Individual rationality

In auctions, we often assume that a buyer is individually rational (IR) and does not expect to lose money. That is, the expected utility $\mathbb{E}[v_i - p_i]$ of a buyer is non negative. The expectation is taken over the distributions of the buyers and for randomized auctions, also over the randomness in the auction's payment rule. (So far, our examples of auctions are all deterministic.)

Individual rationality

In auctions, we often assume that a buyer is individually rational (IR) and does not expect to lose money. That is, the expected utility $\mathbb{E}[v_i - p_i]$ of a buyer is non negative. The expectation is taken over the distributions of the buyers and for randomized auctions, also over the randomness in the auction's payment rule. (So far, our examples of auctions are all deterministic.)

Given that the payment is an expected payment, there are three types of IR that should be distinguished. In what follows we are still assuming common knowledge of everyones distributions and bidding strategies but not assuming i.i.d or any other assumptions about the distributions. Following the notation in KP, we let $b_i(v_i) = \beta(v_i)$ be the bid when the buyer knows his value v_i and has common knowledge about the other buyers.

Three types if IR

The following distinctions for individual rationality also apply to incentive compatability (i.e. truthfulness) of a mechanism that we will discuss later.

- *ex-ante IR*
  $\mathbb{E}[u_i[V_i|V_{-i}] \geq 0]$; that is, the expectation, is over everyones distrbutions including player i.
- *ex-interim IR*
  $\mathbb{E}[u_i[b_i(v_i)|V_{-i}] \geq 0]$; that is, the expectation, is over the other buyers' distrbutions but knowing ones own valuation.
- *ex-post IR*
  $u_i(b_i(v_i)|\mathbf{b}_i) \geq 0$; that is, when the auction is finished and everones true values are revealed, buyer i is guaranteed to have non-negative utility.

The Revelation Principle: a simple but useful insight

We have just touched on a few of the many types of auctions one can have even when there is only one item for sale. We know that some auctions may not be truthful and trying to determine good bidding strategy in a Bayes-Nash equilibirum may be complicated.

The notion of truthfulness we want is called Bayes-Nash incentive compatability (BIC) which we will define to mean that given ones true value, and distributional knowledge of everyone elses strategies, it is a dominant strategy (in expectation) to be truthful.

The Revelation Principle: a simple but useful insight

We have just touched on a few of the many types of auctions one can have even when there is only one item for sale. We know that some auctions may not be truthful and trying to determine good bidding strategy in a Bayes-Nash equilibirum may be complicated.

The notion of truthfulness we want is called Bayes-Nash incentive compatability (BIC) which we will define to mean that given ones true value, and distributional knowledge of everyone elses strategies, it is a dominant strategy (in expectation) to be truthful.

The following is a simple but useful observation. It basically says that any complication of ones bidding strategy in an auction can be simulated by a more complicated auction. (Note : The text assumes that the initial auction is a direct auction but with care it can be extended.)

The revelation principle

If $\mathcal{A}$ is an auction achieving a Bayes-Nash euilibrium, then there is a direct auction (i.e. a sealed bid auction) $\mathcal{A}'$ which is BIC and has the same set of winners and payments as the auction $\mathcal{A}$.

Reserve prices and the Myerson optimal auction for one item

As stated before, the seller may value the item more than the buyers and the way to avoid the seller being disadvantaged is to have a reserve price. We have also already noted that revenue equivalence holds even if we have a reserve price. Even if the auctioneer has little or no value for the item, she can still use a reserve price if that will increase her expected revenue.

Reserve prices and the Myerson optimal auction for one item

As stated before, the seller may value the item more than the buyers and the way to avoid the seller being disadvantaged is to have a reserve price. We have also already noted that revenue equivalence holds even if we have a reserve price. Even if the auctioneer has little or no value for the item, she can still use a reserve price if that will increase her expected revenue.

When should a seller set a reserve price above her actual valuation?

Reserve prices and the Myerson optimal auction for one item

As stated before, the seller may value the item more than the buyers and the way to avoid the seller being disadvantaged is to have a reserve price. We have also already noted that revenue equivalence holds even if we have a reserve price. Even if the auctioneer has little or no value for the item, she can still use a reserve price if that will increase her expected revenue.

When should a seller set a reserve price above her actual valuation? This depends on the tradeoff betwen the probability of an increased sale price and the probability that the item doesn't sell.

Reserve prices and the Myerson optimal auction for one item

As stated before, the seller may value the item more than the buyers and the way to avoid the seller being disadvantaged is to have a reserve price. We have also already noted that revenue equivalence holds even if we have a reserve price. Even if the auctioneer has little or no value for the item, she can still use a reserve price if that will increase her expected revenue.

When should a seller set a reserve price above her actual valuation? This depends on the tradeoff betwen the probability of an increased sale price and the probability that the item doesn't sell.

So is there a way to optimize one's revenue in a single item auction?

Reserve prices and the Myerson optimal auction for one item

As stated before, the seller may value the item more than the buyers and the way to avoid the seller being disadvantaged is to have a reserve price. We have also already noted that revenue equivalence holds even if we have a reserve price. Even if the auctioneer has little or no value for the item, she can still use a reserve price if that will increase her expected revenue.

When should a seller set a reserve price above her actual valuation? This depends on the tradeoff betwen the probability of an increased sale price and the probability that the item doesn't sell.

So is there a way to optimize one's revenue in a single item auction?

Surprinsgly, for independent buyer distributions that are *regular*, there is an optimal auction where the seller just needs to post an appropriate reserve price and then use a 2nd price auction.

Virtual values

We well state the Myerson optimal auction (Definition 14.9.10) and theorem (Theorem 14.9.11) for one item and n buyers with independent (not necesssarily identical) distributions $\{F_i\}$. We will then go back to discuss some cases and the proof. First we need a definition.

The virtual value

Let F be a cumulative distibution function and f is density function (i.e. the derivative of the CDF F). The virtual value function for an agent with distribution CDF F_i is defined as:

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

F_i is *regular* if its virtual value function is strictly increasing.

Claim: common distributions are regular.

The Myerson auction

The Myerson auction

- Agents submit bids b_i
- The item is allocated to the bidder i with the largest virtual value $\psi_i(b_i) \geq 0$; otherwise the item is not allocated.
- The winning buyer (if any) is charged his *threshold bid* $t_*(\mathbf{b}_i) = \min\{b : \psi_i(b) \geq \max(0, \{\psi_j(b_j)\}_{j \neq i)}\}$

This then is a second price auction using virtual bids rather than the actual bids.

Note then that given that the $\{F_i\}$ are not necessarily identical, we can have $\psi_i(v_i) > \psi_j(v_j)$ with $v_j > v_i$ so that this is not a standard auction. (We will see that this is a truthful mechanism so we can assume that $b_i = v_i$ for each buyer.) It follows that for independent (but not identical) distributions, social welfare may not be optimized.

The Myerson theorem

Myerson theorem

If all buyer distributions are regular and independent, then the Myerson auction is optimal for the seller and (ex-post) IC and IR for the buyers.

The fact that truthfulness is a dominant strategy (no matter what the other buyers bid) follows as in the Vickery auction since the Myerson auction is a Vickery auction on the virtual bids.

Corollary for i.i.d. distributions

For buyers with a regular i.i.d. distribution F, the Myerson auction is the Vickery auction with virtual bids and reserve price $\psi^{-1}(0)$.

Myerson auction: a simple example

Lets consider a simple i.i.d. example where the distribution is the uniform distribution $U[0, 1]$.

Myerson auction: a simple example

Lets consider a simple i.i.d. example where the distribution is the uniform distribution $U[0,1]$.

In this case, $F(v) = v$ and $f(v) = 1$. Thus the virtual valuation $\psi(v) = v - \frac{1-F(v)}{1} = 2v - 1$ so that $\psi^{-1}(0) = \frac{1}{2}$.

Myerson auction: a simple example

Lets consider a simple i.i.d. example where the distribution is the uniform distribution $U[0,1]$.

In this case, $F(v) = v$ and $f(v) = 1$. Thus the virtual valuation $\psi(v) = v - \frac{1-F(v)}{1} = 2v - 1$ so that $\psi^{-1}(0) = \frac{1}{2}$.

The Myerson auction then sets a reserve price of $\frac{1}{2}$ and alloactes to the buyer whose virtual bid = virtual value (assuming truthful bidding which is a dominant strategy) is highest (and above the reserve price). .

Myerson auction: a simple example

Lets consider a simple i.i.d. example where the distribution is the uniform distribution $U[0,1]$.

In this case, $F(v) = v$ and $f(v) = 1$. Thus the virtual valuation $\psi(v) = v - \frac{1-F(v)}{1} = 2v - 1$ so that $\psi^{-1}(0) = \frac{1}{2}$.

The Myerson auction then sets a reserve price of $\frac{1}{2}$ and alloactes to the buyer whose virtual bid = virtual value (assuming truthful bidding which is a dominant strategy) is highest (and above the reserve price). .

For two buyers, you can calculate the expected revenue and see that it is greater than the expected revenue of the 1st and 2nd price auctions for which we know that both have expected revenue $= \frac{1}{3}$.

Should we be surprised by having such a simple optimal mechanism?

Should we be surprised by having such a simple optimal mechanism?

Surprise is, of course, in the eyes of the beholder. For selling a single item to one buyer, we can more easily derive what the seller should do and obviously it will be consistent with the Myerson auction.

Should we be surprised by having such a simple optimal mechanism?

Surprise is, of course, in the eyes of the beholder. For selling a single item to one buyer, we can more easily derive what the seller should do and obviously it will be consistent with the Myerson auction.

And the revelation principle does tells us that we can focus on sealed bid (i.e. direct) mechanisms. But there are many types of direct bid mechanisms even for one buyer.

Should we be surprised by having such a simple optimal mechanism?

Surprise is, of course, in the eyes of the beholder. For selling a single item to one buyer, we can more easily derive what the seller should do and obviously it will be consistent with the Myerson auction.

And the revelation principle does tells us that we can focus on sealed bid (i.e. direct) mechanisms. But there are many types of direct bid mechanisms even for one buyer.

In general (for one item), the auctioneer can offer a *menu* set of $\{Prob_i[winning], price_i\}$. So to me it is surprising that (as long as the seller knows the distributions), then a simple reserve price auction is optimal (for the seller).

Should we be surprised by having such a simple optimal mechanism?

Surprise is, of course, in the eyes of the beholder. For selling a single item to one buyer, we can more easily derive what the seller should do and obviously it will be consistent with the Myerson auction.

And the revelation principle does tells us that we can focus on sealed bid (i.e. direct) mechanisms. But there are many types of direct bid mechanisms even for one buyer.

In general (for one item), the auctioneer can offer a *menu* set of $\{Prob_i[winning], price_i\}$. So to me it is surprising that (as long as the seller knows the distributions), then a simple reserve price auction is optimal (for the seller).

And as section 14.11.1 shows, the situation becomes much more "interesting" as soon as we have two items, even with just one additive buyer having an additive valuation function and one seller.