CSC304: Algorithmic Game Theory and Mechanism Design Fall 2016

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

October 17, 2016

Lecture 10

- Announcements
  I have posted the first two questions for Assignment 2 which is due October 28. I expect to post some more questions during the week.
- Todays agenda
  We continue the discussion of auctions.
  - We continue the discussion of the Vickery 2nd price mechanism for selling a single item.
  - We compare the Vickery auction and the 1st price auction.
  - We will use today to give a very fast overview of some of the important results in Chapter 14 of the KP text. 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.

The English (ascending price) auction

While the second price auction may seem a little strange at first, it is in some sense equivalent to the familar English ascending price auction.

The English (ascending price) auction

While the second price auction may seem a little strange at first, it is in some sense equivalent to the familar English ascending price auction.

In the English auction, the auctioneer starts with a price (i.e. a reserve price) and then continues to ask who wants to raise the bid? The final (highest) bid wins.

The English (ascending price) auction

While the second price auction may seem a little strange at first, it is in some sense equivalent to the familar English ascending price auction.

In the English auction, the auctioneer starts with a price (i.e. a reserve price) and then continues to ask who wants to raise the bid? The final (highest) bid wins.

Suppose the auctioneer always asks if there is anyone who wants to raise the current bid by some small ϵ. Then this ascending price auction is essentially producing the same outcome (assuming buyers do not change their valuation given other bids) in terms of who wins the item and what is paid. These are referred to as two different *implementations* of the same outcome.

The English (ascending price) auction

While the second price auction may seem a little strange at first, it is in some sense equivalent to the familar English ascending price auction.

In the English auction, the auctioneer starts with a price (i.e. a reserve price) and then continues to ask who wants to raise the bid? The final (highest) bid wins.

Suppose the auctioneer always asks if there is anyone who wants to raise the current bid by some small ϵ. Then this ascending price auction is essentially producing the same outcome (assuming buyers do not change their valuation given other bids) in terms of who wins the item and what is paid. These are referred to as two different *implementations* of the same outcome.

But how do these two implementations differ?

The English (ascending price) auction

While the second price auction may seem a little strange at first, it is in some sense equivalent to the familar English ascending price auction.

In the English auction, the auctioneer starts with a price (i.e. a reserve price) and then continues to ask who wants to raise the bid? The final (highest) bid wins.

Suppose the auctioneer always asks if there is anyone who wants to raise the current bid by some small ϵ. Then this ascending price auction is essentially producing the same outcome (assuming buyers do not change their valuation given other bids) in terms of who wins the item and what is paid. These are referred to as two different *implementations* of the same outcome.

But how do these two implementations differ?

The Dutch descending price auction (i.e. keep decreasing the price by a small amount until one buyer accepts the price) is in the same way "equivalent" to a 1st price auction .

Revenue equivalence

We should have seen (if my iphone 4 sale worked as it should have) that the buyers strategies are different for 1st price and 2nd price auctions. Namely, if you believed that the 2nd price auction was truthful, then you would bid your true value whereas for a 1st price auction you should have "shaved" your bid as otherwise your utility is guaranteed to be 0 whether you win or lose the auction. We can discuss why it didn't turn out as the theory would imply. For the 2nd price (resp. 1st price) auction, the average bid/value ratio was .632 (resp. .721).

Which mechanism is best for me (the seller) in terms of my revenue?

Revenue equivalence

We should have seen (if my iphone 4 sale worked as it should have) that the buyers strategies are different for 1st price and 2nd price auctions. Namely, if you believed that the 2nd price auction was truthful, then you would bid your true value whereas for a 1st price auction you should have "shaved" your bid as otherwise your utility is guaranteed to be 0 whether you win or lose the auction. We can discuss why it didn't turn out as the theory would imply. For the 2nd price (resp. 1st price) auction, the average bid/value ratio was .632 (resp. .721).

Which mechanism is best for me (the seller) in terms of my revenue?

Surpising fact: If all bidders are independently drawing their values from same uniform distribution $U[0, h]$ then the expected revenue (in equilibrium) is the same for the first and second price auctions!

Revenue equivalence

We should have seen (if my iphone 4 sale worked as it should have) that the buyers strategies are different for 1st price and 2nd price auctions. Namely, if you believed that the 2nd price auction was truthful, then you would bid your true value whereas for a 1st price auction you should have "shaved" your bid as otherwise your utility is guaranteed to be 0 whether you win or lose the auction. We can discuss why it didn't turn out as the theory would imply. For the 2nd price (resp. 1st price) auction, the average bid/value ratio was .632 (resp. .721).

Which mechanism is best for me (the seller) in terms of my revenue?

Surpising fact: If all bidders are independently drawing their values from same uniform distribution $U[0, h]$ then the expected revenue (in equilibrium) is the same for the first and second price auctions!

Even more surprising: This revenue equivalence holds more generally given some technical (but intuitive) conditions.

The revenue equivalence for 1^{st} and 2^{nd} price auction and the uniform distribution $U[0,1]$.

This was discussed in tutorial so lets just do this quickly and informally.

- In a second price equilibrium, we can assume buyers bid truthfully and pay the second (highest) bid (= value). For the uniform distribution $U[0,1]$, we claim that the expectation of the k^{th} lowest value is $\frac{k}{n+1}$. The expectation of the second highest value is then $\frac{n-1}{n+1}$. This implies that the expected revenue is $\frac{n-1}{n+1}$.
- For a first price auction, the buyer will bid (at equilibrium) a fraction $\frac{n-1}{n}$ of his value and the highest value has expectation $\frac{n}{n+1}$ so that the expected bid (and price) for the winning buyer is $\frac{n-1}{n}\frac{n}{n+1} = \frac{n-1}{n+1}$.

Revenue Equivalence Theorem

Here 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 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.

What do all the technical conditions mean?

Revenue equivalence theorem continued

The theorem assumes all bidders are drawing their value from the same continuous distribution. It isn't said explicitly but like the presentation in the KP text, the proof uses the fact that the payment rule and the allocation probability are differentiable (as well as increasing) functions of one's value.

Revenue equivalence theorem continued

The theorem assumes all bidders are drawing their value from the same continuous distribution. It isn't said explicitly but like the presentation in the KP text, the proof uses the fact that the payment rule and the allocation probability are differentiable (as well as increasing) functions of one's value.

Since every buyer is drawing from the same distribution, it is reasonable to assume that in an equilibrium, all players will be playing the same strategy so that we are considering a symmetric equilbrium.

Revenue equivalence theorem continued

The theorem assumes all bidders are drawing their value from the same continuous distribution. It isn't said explicitly but like the presentation in the KP text, the proof uses the fact that the payment rule and the allocation probability are differentiable (as well as increasing) functions of one's value.

Since every buyer is drawing from the same distribution, it is reasonable to assume that in an equilibrium, all players will be playing the same strategy so that we are considering a symmetric equilbrium.

A *standard auction* is one that will always allocate the item to a player with the highest bid.

Revenue equivalence theorem continued

The theorem assumes all bidders are drawing their value from the same continuous distribution. It isn't said explicitly but like the presentation in the KP text, the proof uses the fact that the payment rule and the allocation probability are differentiable (as well as increasing) functions of one's value.

Since every buyer is drawing from the same distribution, it is reasonable to assume that in an equilibrium, all players will be playing the same strategy so that we are considering a symmetric equilbrium.

A *standard auction* is one that will always allocate the item to a player with the highest bid.

Intuitively, A risk neutral buyer is simply trying to maximize their utility without taking into consideration either minimizing the probabiity of low utility (i.e. a risk averse buyer) or maximizing the probabilty of high utility (i.e. a risk seeking buyer).

The revenue theorem continued

A buyer with value 0 for the item is unlikely want to offer any payment as that would lead to negative utility. More generally, we just need some normalizing condition; e.g. if the distribution has suppport $[h_1, h_2]$ then the expected payment $p(h_1) = p'$ for some fixed p' independent of the particular auction.

The revenue theorem continued

A buyer with value 0 for the item is unlikely want to offer any payment as that would lead to negative utility. More generally, we just need some normalizing condition; e.g. if the distribution has suppport $[h_1, h_2]$ then the expected payment $p(h_1) = p'$ for some fixed p' independent of the particular auction.

We emphasize that there is a wide range of possible auctions that fit these conditions, including the 1st and 2nd (i.e. Vickery) price auctions as well as say a third price auction, the all pay auction and many others.

The revenue theorem continued

A buyer with value 0 for the item is unlikely want to offer any payment as that would lead to negative utility. More generally, we just need some normalizing condition; e.g. if the distribution has suppport $[h_1, h_2]$ then the expected payment $p(h_1) = p'$ for some fixed p' independent of the particular auction.

We emphasize that there is a wide range of possible auctions that fit these conditions, including the 1st and 2nd (i.e. Vickery) price auctions as well as say a third price auction, the all pay auction and many others.

The all pay auction allocates to the highest bidder but charges all buyers their bid. As the text notes, this is an auction implicitly used when bidding for contracts where there is a cost in making the bid.

The revenue theorem continued

A buyer with value 0 for the item is unlikely want to offer any payment as that would lead to negative utility. More generally, we just need some normalizing condition; e.g. if the distribution has suppport $[h_1, h_2]$ then the expected payment $p(h_1) = p'$ for some fixed p' independent of the particular auction.

We emphasize that there is a wide range of possible auctions that fit these conditions, including the 1st and 2nd (i.e. Vickery) price auctions as well as say a third price auction, the all pay auction and many others.

The all pay auction allocates to the highest bidder but charges all buyers their bid. As the text notes, this is an auction implicitly used when bidding for contracts where there is a cost in making the bid.

We also emphasize that there is a wide range of possible valuation distributions although to provide simple examples, one often considers uniform distributions $U[h_1, h_2]$ and more specifically $U[0, 1]$.

Sketch of proof (sections 14.4.1 and 14.4.2)

We assume that the bidders are drawing their values i.i.d from a distribution with strictly increasing cummulative distribution function F. We are considering a symmetric equilibrium for bidding $\beta = \beta_i$ for all i. We need to define the probability $a_i(v)$ that the buyer i wins the auction when bidding $\beta(v_i)$ and then paying $p_i(v_i)$ assuming that all the other players j are playing the same strategy $\beta(v_j)$. We will drop the subscript on the probability a and payment p since all buyers are playing the same strategy.

Sketch of proof (sections 14.4.1 and 14.4.2)

We assume that the bidders are drawing their values i.i.d from a distribution with strictly increasing cummulative distribution function F. We are considering a symmetric equilibrium for bidding $\beta = \beta_i$ for all i. We need to define the probability $a_i(v)$ that the buyer i wins the auction when bidding $\beta(v_i)$ and then paying $p_i(v_i)$ assuming that all the other players j are playing the same strategy $\beta(v_j)$. We will drop the subscript on the probability a and payment p since all buyers are playing the same strategy.

The proof comes in two parts, where the first part (i) proves a necessary condition for β to be an equilibrium and the resulting probability $a(v_i)$ and payment $p(v_i)$. The second part verifies that this does indeeed yield an equilibrium.

Sketch of proof (sections 14.4.1 and 14.4.2)

We assume that the bidders are drawing their values i.i.d from a distribution with strictly increasing cummulative distribution function F. We are considering a symmetric equilibrium for bidding $\beta = \beta_i$ for all i. We need to define the probability $a_i(v)$ that the buyer i wins the auction when bidding $\beta(v_i)$ and then paying $p_i(v_i)$ assuming that all the other players j are playing the same strategy $\beta(v_j)$. We will drop the subscript on the probability a and payment p since all buyers are playing the same strategy.

The proof comes in two parts, where the first part (i) proves a necessary condition for β to be an equilibrium and the resulting probability $a(v_i)$ and payment $p(v_i)$. The second part verifies that this does indeeed yield an equilibrium.

Necessary condition for a symmetric equilibrium

Since we are assuming a standard auction where the highest bid wins, and F is strictly increasing, it is reasonable that $\beta(v)$ is also strictly increasing and then any bidding strategy is equivalent to bidding $\beta(w)$ for some w when given v_i.

Necessary condition for a symmetric equilibrium

Since we are assuming a standard auction where the highest bid wins, and F is strictly increasing, it is reasonable that $\beta(v)$ is also strictly increasing and then any bidding strategy is equivalent to bidding $\beta(w)$ for some w when given v_i.

We have

$$a(w) = Prob[\beta(w) \max_{j \neq i} \beta(V_j)] = Prob[\beta(w) \max_{j \neq i} V_j] = F^{n-1}$$

which gives utility $u(w|v_i) = v_i a(w) - p(w)$

Necessary condition for a symmetric equilibrium

Since we are assuming a standard auction where the highest bid wins, and F is strictly increasing, it is reasonable that $\beta(v)$ is also strictly increasing and then any bidding strategy is equivalent to bidding $\beta(w)$ for some w when given v_i.

We have

$$a(w) = Prob[\beta(w) \max_{j \neq i} \beta(V_j)] = Prob[\beta(w) \max_{j \neq i} V_j] = F^{n-1}$$

which gives utility $u(w|v_i) = v_i a(w) - p(w)$

For $\beta(v_i)$ to be an equilibrium, the derivative $v_i a'(w) - p'(w)$ of the utility must be 0 at $w = v_i$. That is, $p'(v_i) = v_i a'(v_i)$ for all v_i.

Necessary condition for a symmetric equilibrium

Since we are assuming a standard auction where the highest bid wins, and F is strictly increasing, it is reasonable that $\beta(v)$ is also strictly increasing and then any bidding strategy is equivalent to bidding $\beta(w)$ for some w when given v_i.

We have

$$a(w) = Prob[\beta(w) \max_{j \neq i} \beta(V_j)] = Prob[\beta(w) \max_{j \neq i} V_j] = F^{n-1}$$

which gives utility $u(w|v_i) = v_i a(w) - p(w)$

For $\beta(v_i)$ to be an equilibrium, the derivative $v_i a'(w) - p'(w)$ of the utility must be 0 at $w = v_i$. That is, $p'(v_i) = v_i a'(v_i)$ for all v_i.

Using (for simplicity), the particular normalization $p(0) = 0$, and integrating, we get $p(v_i) = \int_0^{v_i} v \cdot a'(v) dv$
so that integrating parts, we get $p(v_i) = v_i a(v_i) - \int_0^{v_i} a(w) dw$

Completing the proof of revenue equivalence

Note that since $a(v) = F^{n-1}$, this probability only depends on the distribution and does not depend on the auction as long as the winner is the buyer with the highest valuation which will be the case since this is a symmetric auction and the highest bid corresponds to the highest value). Furthermore, since $p(v)$ only depends on v and $a(v)$, the price is also just a function of the distribution. Finally, the price $p(v)$ for each buyer determines the expected revenue from each buyer which determines the total expected revenue. Hence we obtain revenue equivalence,

Completing the proof of revenue equivalence

Note that since $a(v) = F^{n-1}$, this probability only depends on the distribution and does not depend on the auction as long as the winner is the buyer with the highest valuation which will be the case since this is a symmetric auction and the highest bid corresponds to the highest value). Furthermore, since $p(v)$ only depends on v and $a(v)$, the price is also just a function of the distribution. Finally, the price $p(v)$ for each buyer determines the expected revenue from each buyer which determines the total expected revenue. Hence we obtain revenue equivalence,

It only remains to verify that the bidding strategy $\beta(v)$ is an equilibrium which is done in section 14.4.2. That is, the proof shows that the utility of bidding $\beta(v)$ given value v dominates bidding $\beta(w)$ for any $w \neq v$.

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.

- I believe (but have to verify) that it can be extended to i.i.d. discrete distributions although in the Economics literature, it seems that one is always considering continuous distributions.
- 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.