Lecture Notes. *Expected Revenue in Auctions.*

So far we have analyzed each auction as a game and found the equilibrium outcome for each auction format. To do that we took one of the buyers and found his optimal bid (best response to other buyers' bidding strategies). We expressed optimal bid as a function of a buyer's own valuation V_i. Then we used optimal bids to find the equilibrium bids and the outcome (the winner and the price) for each auction. In the beginning of this set of lectures we discussed that a good auction if efficient and revenue maximizing. We found that in all auctions the winner is the buyer who values the object the most, therefore, all auction formats are efficient. How about revenue maximization? From the information we gathered so far it is not clear whether first-price or second-price auctions raise higher revenue, so we would not be able to give any recommendations as to which auction format to choose to get highest possible price. We are going to fill that gap.

It should be more or less obvious that if the auctioneer knows the ***actual*** valuations of the buyers (while the buyers don't know each other's V's), he could possibly choose between second- and first-price auctions. The auctioneer can take optimal bidding strategies and calculate the price that he will get if he conducts a first or a second- price auction. But, of course, most of the time actual valuations are not observable to auctioneer, who at best would know the distribution of the valuations. *So far we reached the conclusion that without information about the actual valuations of the buyers who participate in an auction, there is no reason to expect that any particular auction format would raise a higher revenue.* Recall that when someone is faced with uncertainty, he tries to choose an action that maximizes his expected payoff. Without knowing the valuations of the buyers, the best an auctioneer can do is to choose an auction format that maximizes *expected revenue*. Let's see what is the expected revenue raised by all auctions and compare it. At first we will look at a general case, without making any assumptions about the distribution of the valuations, then we will check the revenue equivalence for the uniform distribution.

Expected revenue in English and second-price sealed-bid auctions.

Recall that the outcome of these two auctions is the same: the winner is the buyer with the highest valuation and the *price is equal to the second-highest* ***actual*** *valuation*. The expected outcome is easy to describe in terms of expected valuations:

- The *expected revenue is equal to the expected second-highest valuation*, given the number of buyers and the distribution of the valuations. Let's denote expected revenue of second-price auctions as ER^{2ND}:

$$ER^{2ND} = \textit{Expected } 2^{nd}\textit{-highest valuation.}$$

Expected revenue in Dutch and first-price sealed-bid auctions.

 [idudnyk@sfu.ca]

So far we know that in first-price auctions optimal bids are always below the true valuations of the buyers; the winner will pay some fraction of his (highest, of course,) valuation. Therefore, the expected revenue of a first-price auction is:

$$ER^{1ST} = \textit{fraction of the highest valuation.}$$

How can we compare ER^{1ST} and ER^{2ND}? In order to be able to compare the expected revenue in different auctions we need to take a closer look at the optimal bid in the first-price auctions. We will see, that if the buyer with the highest valuation is smart (uses all available information to maximize his expected utility), in a first-price auction, on average, he will not pay more that he pays in a second-price auction.

Perfect Information. As an exercise let's at first think about what would be the outcome of a first-price auction of all buyers knew everything they need to know to calculate their optimal bids. ***IF*** *buyers had all necessary information (could observe each other's valuations) what would be the optimal bids and the outcome of a first-price auction?*

What is the NE of a first-price auction if valuations are observable? IF the buyer with the highest valuation could see valuations of all other buyers, he would place bid equal to the second-highest valuation and win the object at that price. Surprising, but true, *if there is no private information, first-price auctions raise* ***exactly the same revenue*** *as the second-price auctions.* In conclusion, *if valuations are observable, the optimal bid of the buyer with the highest valuation is equal to the second-highest valuation.*

Under **asymmetric information** buyers cannot observe each other's valuations. Therefore, each buyer is wondering whether his valuation is the highest and if it is, what is the second-highest valuation (the valuation of the buyer whom he needs to outbid). How should a buyer find his optimal bid without knowing those things? Well, he can *use the same logic, but instead of actual values,* he will have to *use expected values.* Let's pretend we are the buyer indexed i and try to find the optimal bid only knowing own valuation V_i, N, and distribution of valuations. The 'buyer i should be thinking the following:

- **Am I going to win?** Don't know. I only know that: *if I win then my own valuation, V_i* ***must be*** *the highest.*
  The best I can do at this point is to ***assume*** *that I am going to win.* Therefore, SUPPOSE, my valuation V_i is the highest.
- **How much should I bid?** In order to win I need to outbid the buyer with the 2^{nd}-highest valuation. I don't know the actual valuation I need to outbid, what should I do? Calculate the expected value. If I win, my valuation is the highest, I can calculate expected 2^{nd}-highest valuation using the number of buyers and the distribution of valuations, ASSUMING that V_i is the highest valuation. My ***optimal bid*** *is equal to the expected second-highest valuation* ***assuming*** *that own valuation V_i is the highest (as you will see in a later example, this is not exactly the same as simply bidding expected second-highest valuation).*

Since every buyer will follow the same reasoning, the winner will be the buyer with the highest valuation (because his estimate of 2^{nd} highest valuation will be the highest) and the price will be equal to his bid. Now we can describe the revenue in terms of expected valuations:

- In a first-price auction the buyer with the highest valuation wins by bidding his estimate of the second-highest valuation assuming that own valuation is the highest. 'On average' winner will have valuation equal to the expected highest value if N values are drawn from the distribution, therefore, the expected winning bid on average will be equal to the expected 2^{nd}-highest valuation:

$$ER^{1ST} = \textit{Expected } 2^{nd}\textit{-highest valuation.}$$

Revenue Equivalence Theorem.

Let's do a short summary. If information is symmetric, the auctioneer will get actual revenue $P = 2^{nd}$-highest actual valuation, no matter what auction format he chooses. If information is asymmetric, we have to operate with expected revenue and expected valuations. We have shown intuitively that in a first-price auction, essentially, the buyer with the highest valuation will optimally bid his estimate of the 2^{nd}-highest valuation. Therefore, *all four standard auction formats raise the same expected revenue.* The formal mathematical proof of this claim is beyond the scope of this course, but mathematical formulas essentially tell the same story about calculating optimal bids in first-price auctions.

Let's formally state the ***revenue equivalence theorem***. Any theorem tells you that if some conditions are satisfied, some result must be true. The revenue equivalence theorem tells you that:

IF:

- the number of buyers, N is fixed (the same)
- buyers' valuations are private and independent
- valuations are drawn from the same distribution
- buyers' utility is linear in wealth[1], $U_i = V_i - P$, (which is what we are using for this class)
- the buyer with the lowest valuation gets the same utility in all auctions (there is no/same fixed entry fee or reservation price/bid in all auctions)

THEN *all four standard auction formats raise the same expected revenue.*

[1] This assumption is related to preferences towards risk, which is not covered in this course and will not appear on the exams in any form. Linear utility means that buyers are risk-neutral, they are indifferent between having 100 dollars for sure or facing a lottery that gives 0 or 200 with equal probability, because risk-neutral people only care about expected values. If buyers are risk averse (don't like to face uncertainty) the revenue equivalence does not hold. Bids will not change in the 2^{nd}-price auctions (dominant strategy), but in a first-price auction buyers will bid high proportion of their valuation to minimize the risk of losing the object. In conclusion if buyers are risk averse first-price auctions will raise higher expected revenue

Fixed number of buyers, same distribution, and same utility of the buyer with the lowest valuation is needed to make sure you are not comparing pears to oranges. Essentially this theorem tells you that if you 'invite' a fixed number of buyers who's valuations follow some particular distribution, *ex-ante*, before the bidding starts, it does not matter which auction you conduct, because all raise exactly the same ***expected*** revenue and without knowing anything about the actual valuations, there is no reason for you to choose one auction format over the other. *Ex-post*, or after the bids are submitted (and you can see/guess what were the ***actual*** valuations were), you might realize that you made a mistake and other auction format would raise a higher price, but after the bids are submitted, there is not much you can do about that.

Does Revenue Equivalence work with the uniform distribution?

Let's check the revenue equivalence for the uniform distribution. At fist, we will show that when the valuations are uniformly distributed between zero and one all auctions raise the same expected revenue; to do that we will only use the optimal bidding strategies we derived (true bidding in the second price auctions and fraction-of-the-true-valuation bidding in the first-price auctions). Then, to make sure that I did not give you any inconsistent information, we will show that in the symmetric BNE that we have found for the first-price auctions, buyers were actually bidding what they expect is the second highest valuation, assuming that own valuation is the highest. Since we will be dealing with math and formulas, let's first see what the expected highest and second-highest valuation look like if N valuations are drawn from uniform distribution between zero and one.

Expected valuations drawn from uniform distribution. If N valuations are drawn randomly from a uniform distribution between zero and one, the expected i^{th}-highest valuation is:

$$\frac{N-i+1}{N+1}$$

The expected highest valuation, $i=1$, is:

$$\frac{N}{N+1}$$

the expected 2^{nd}-highest valuation, $i=2$, is:

$$\frac{N-1}{N+1}$$

the expected lowest valuation, $i=N$ is:

$$\frac{1}{N+1}$$

Now you are equipped with the formulas to calculate the expected revenue in second- and first- price auctions. **Comment:** expected values if N outcomes are drawn from some distribution are called "order statistics". The formulas are standard and can be found in statistics textbooks. Of course, these formulas will be different for different distributions. For this course you only need to know how to operate with the formulas for uniform distribution.

2^{nd}***-price auctions***. If there are N buyers, the expected revenue raised by English and 2^{nd}-price sealed-bid auctions is equal to the expected second-highest valuation:

$$ER^{2ND} = \frac{N-1}{N+1}$$

1^{st}***-price auctions***. In Dutch and first-price sealed-bid auctions, the expected revenue is equal to the expected winning bid. We found that in BNE all buyers, including the winner, bid $\lambda = \frac{N-1}{N}$ fraction of their true valuation. The expected revenue is equal to λ fraction of the expected highest valuation :

$$ER^{1ST} = \lambda \frac{N}{N+1} = \frac{N-1}{N} \cdot \frac{N}{N+1} = \frac{N-1}{N+1}$$

Both auctions raise the same expected revenue equal to $\frac{N-1}{N+1}$, where N is the number of buyers. The expected revenue is equal to the expected second-highest valuation.

Finally, in the symmetric BNE, do buyers optimally bid 'expected second-highest valuation, assuming own valuation is the highest?'. We need one more formula. If buyer i assumes that his own valuation V_i is the highest, what is his expectation about the second-highest valuation, which is supposed to be his optimal bid in a general case? If some buyer's valuation V_i is the highest, meaning that the remaining $N-1$ valuations are uniformly distributed on $[0, V_i]$ The expected second-highest (the highest of the $N-1$ 'other buyers' ' valuations) takes value of

$$\frac{N-1}{N} V_i$$

This is exactly the bid submitted by a buyer with valuation V_i in the symmetric Bayesian Nash Equilibrium. Notice that we found that the optimal bid $b_i^* = \frac{N-1}{N} V_i$ by simple maximization of expected utility, we did not use any understanding of what conditions the optimal bid should satisfy.

Practice: Suppose that the seller knows that there are 4 buyers, the actual highest valuation is equal to .8. Which auction format do you recommend? *(Hint: there is a trick. Now that the seller knows the highest valuation, he can use that information for calculating the second-highest valuation.)*

[idudnyk@sfu.ca]