Econ 805 – Advanced Micro Theory I
Dan Quint
Fall 2009
Lecture 16

Today - a fairly brief overview of multi-unit auctions.

- Everything we've done so far has assumed unit demand. In a couple of cases, we've allowed for the possibility of multiple objects – either sold by the same seller, or by distinct sellers – but we've continued to assume that each buyer only wants a single object.
- Once we allow for bidders wanting multiple auctions, things immediately get much more complicated.
- The Milgrom book divides auctions in this setting into two cases, depending on whether they follow the *law of one price*, that is, whether identical objects sell for the same price

Uniform Price Sealed-Bid Auctions

The simplest auction for multiple identical units is a uniform-price sealed-bid auction. Each bid is a price-quantity pair, for example, a bid to buy 5 units at a price of $20 each. Bidders can submit multiple bids, at different prices – say, 5 units at $20, and another 5 units at $15. The auctioneer aggregates all these bids into a single aggregate demand curve, and sets a single market-clearing unit price which all winners pay.

Bids can be seen in one of two ways: either as a series of prices bid on individual objects $(20, 20, 20, 20, 20, 15, 15, 15, 15, 15$ in our example), or as a demand function (demand is 0 if the price is above 20, 5 if the price is between 15 and 20, and 10 if the price is below 15).

As with double-auctions, if there are m objects for sale and the prices bid on the m^{th} and $m+1^{st}$ objects are different, there is a range of prices that clear the market. We typically assume that price is set at the highest rejected bid, as in the second-price auction for a single good.

We generally assume that bidders have weakly decreasing valuations for each successive good – that is, they value the second good less than the first, the third less than the second, and so on. It's not really clear how well this auction works when this assumption is violated. (As we'll see, complementarities among objects lead to problems in many auction formats.)

With uniform-price auctions, the main problem to be aware of is **demand reduction** – that bidders who want more than one object have an incentive to shade their demand downwards. (This can be thought of as bidding less than their valuation for incremental units, or by demanding fewer units than they want at a given price.)

It's easiest to see this in an example. First, note that when a bidder is only interested in a single unit – his value for incremental units beyond the first is 0 – it is a dominant strategy to bid his true valuation for one unit. (Same as in the second-price auction.)

Now suppose there are two bidders, and two objects for sale. The first bidder wants only one object, and his valuation for it is $U[0,1]$. You're the second bidder, and you value the first object at v_1 and the second at v_2, with $1 \geq v_1 \geq v_2 \geq 0$.

Suppose you bid b_1 on the first object and b_2 on the second. Since $b_1 \geq b_2$, b_1 can never set the price, since it can never be the highest rejected bid. Let t_1 be the first bidder's valuation, which is also his bid.

If $b_2 > t_1$, you win both objects, and pay a price of t_1 per unit. If $b_2 < t_1$, you win one object, and pay b_2. So expected revenue can be written as

$$\int_0^{b_2} (v_1 + v_2 - 2t_1)dt_1 + \int_{b_2}^1 (v_1 - b_2)dt_1$$

$$= (v_1 + v_2)b_2 - b_2^2 + v_1(1 - b_2) - b_2(1 - b_2) = v_1 + v_2 b_2 - b_2 = v_1 - b_2(1 - v_2)$$

which is strictly decreasing in b_2 – you're best off bidding 0 for the second unit.

Obviously, this example is extreme – in many settings, you still want to bid for multiple units. However, when bidding for any unit after the first, your bid has an impact both on the probability of winning that marginal unit, and on the expected price you pay for the units you've already bid on – the inframarginal units. So it's always optimal to bid less than your true valuation on each unit but the first.

(This is analogous to why monopolists price above marginal cost. With a single unit for sale in a second-price auction, you never set the price when you win, so your bid only determines whether or not you win, not the price you pay. When you're bidding on multiple units, your i^{th} bid could set the price while you still win $i-1$ objects, so there is an incentive to shade your bid, or to reduce your reported demand.)

In his book, Milgrom formalizes this into a theorem, in a setting with bidders who want two items each. The theorem basically says that it's a dominant strategy to bid your valuation for the first object, but there's never an equilibrium where all bidders bid their full value on both, and therefore (as long as the distribution of bidder valuations has full support) there's no equilibrium in which the allocation is always efficient.

Another problem with uniform-price auctions is the existence of "low-revenue equilibria", that is, equilibria in which the seller's revenue is very low. A simple example: suppose there are N objects for sale, and N bidders, each of whom want up to $k > 1$ item and value each incremental item up to this capacity at \$1. At any price below \$1, there is demand for Nk items, so you would expect competition to push the price to \$1 per unit. However, one equilibrium of the uniform-price auction is for each bidder to bid \$1 for one unit, and 0 for all additional units, leading to total revenue of 0. When the number of objects doesn't match the number of sellers, there are similar low- or no-revenue equilibria involving some bidders not bidding on any objects.

Milgrom also constructs an example of similar, low-revenue equilibria in a setting with divisible goods, so the problem has nothing to do with the discrete nature of the items for sale.

Simultaneous Ascending Auctions

Another type of auction sometimes used (in FCC spectrum auctions, among other settings) is the **simultaneous ascending auction**. The items up for sale are treated distinctly, that is, bidders bid on a particular object, not a generic one. A simultaneous ascending auction takes place in a series of rounds. At the end of each round, the auctioneer announces the "standing high bid" for each object, and a minimum bid for each item in the next round, usually a fixed percentage above the standing high bid. In the next round, bidders submit new bids. (In early simultaneous ascending auctions, bidders could submit any bid above the minimum; in later versions, they were limited to a discrete set of choices, to avoid communication via bids.) If multiple bidders submit the same new high bid on an item, one of them is chosen at random to be the new standing high bidder.

All of the auctions remain open until a round when no new bids are submitted, at which point all of the auctions end and the standing high bidders win at those prices. To prevent bidders from waiting until late in the auction and then jumping in, FCC auctions had an **activity rule**: the number of items for which you were "active" (the number for which you were the standing high bidder at the end of a round, plus the number for which you submitted new bids in the next round) cannot go up over time. That is, a bidder who starts out bidding for k objects is limited to bidding on at most k objects for the rest of the auction, although which ones he bids on can change.

The results on simultaneous ascending auctions are these:

- When the items for sale are all substitutes to all bidders, then
  - A competitive equilibrium exists (that is, market-clearing prices exist)
  - Bidding "straightforwardly" is feasible
  - If all bidders bid straightforwardly, the outcome of the simultaneous ascending auction is approximately equivalent to the competitive equilibrium, and approximately efficient (both up to the size of the bid increment)
- However, when the items for sale are not all substitutes to all bidders, then
  - A competitive equilibrium may not exist
  - Bidding straightforwardly is not feasible
  - Bidders face an "exposure problem"

Formally, let N be the set of bidders, and $L = \{1, 2, \ldots, L\}$ the set of objects for sale. We can describe a subset of L by a vector of 1's and 0's indicating whether each item is in the set. If bidder j acquires the allocation of goods x and pays m, his payoff is $v^j(x) - m$. Bidder j's demand correspondence is

$$D^j(p) = \arg\max_x \{v^j(x) - p \cdot x\}$$

We assume free disposal, so v_j is increasing in the vector x.

There will be some prices at which $D^j(p)$ is a set, that is, where a bidder is indifferent between two bundles; these are nongeneric. We'll focus on the rest of the price space. (That is, when we look at auctions, we'll assume that prices are restricted to those on which each bidder's demand is single-valued.)

Goods are defined to be **substitutes** if wherever D^j is single-valued, the demand for a particular item is weakly increasing in the prices of the other items. Formally, goods are substitutes if

$$\hat{p}_{-l} \geq p_{-l}, \hat{p}_l = p_l \quad \longrightarrow \quad D_l^j(\hat{p}) \geq D_l^j(p)$$

(Some of the literature uses an equivalent condition called "gross substitutes." In usual econ terminology, "gross substitutes" is based on Marshallian demand, and "substitutes" is based on Hicksian demand. However, we're in a quasilinear world, so Marshallian and Hicksian demand are equivalent.)

The key to using demand functions in an auction setting is to realize that at a given point in an auction, bidders face different price vectors, based on whether or not they are the standing high bidder for a given object. That is, let $(b_1, \ldots, b_l)$ be the standing high bids for the various items at some point during the auction. Suppose that the minimum bid is $(1+\epsilon)$ times the standing high bid. In the next round, one bidder still has the possibility of paying b_1 for the first item: the standing high bidder. Everyone else faces a minimum price of $(1+\epsilon)b_1$, that is, everyone else knows that at best, they could win the first item for $(1+\epsilon)b_1$.

Define

- S^j as the set of goods for which j is the standing high bidder
- p^j as bidder j's "personalized prices", $p^j = (p_{S_j}, (1+\epsilon)p_{L-S^j})$
- We say bidder j bids *straightforwardly* if at every point in the auction, $S^j \subset D^j(p^j)$, and in the next round, bidder j bids the minimum on the set of goods $D^j(p^j) - S^j$

That is, bidder j bids straightforwardly if at every round, he bids on the set of goods which are the bundle he would demand at his current personalized price vector. (Again, recall that we're assuming the generic condition that prices are such that demand is single-valued for everyone.)

The importance of items being substitutes is this:

Theorem 1. *Straightforward bidding is always feasible for bidder j if and only if goods are substitutes for bidder j.*

As the auction goes on and my opponents bid on some items, my personalized price for those items goes up. The substitutes condition ensures that, if I'm the standing high bidder for item l and the prices of the other items go up, I still want to win item l at my current bid; that is, at my new personalized prices, item l is still in my demand set.

The "only if" part is perhaps the interesting part. Suppose I see two of the items as complements – say, one is a spectrum license for Seattle, one is a spectrum license for Portland, and I only want one if I can have both. I start out bidding on both licenses; but I might get outbid on one of them and “stuck” with the other. This is referred to as the **exposure problem** – a bidder who sees the items as complements runs the risk of winning an object he no longer wants once he’s been outbid for a different item. The substitutes condition is required to ensure that this can’t happen. (Tomorrow’s lecture will be about how to deal with complements.)

Under substitutes, however, we said that straightforward bidding is always feasible. The Milgrom book gives a theorem that if preferences are substitutes and all bidders bid straightforwardly, the outcome of a simultaneous ascending auction is basically (up to the size of the minimum bid increment) efficient, and the payments basically match the competitive equilibrium (the set of market-clearing prices).

Since we’re relating auction results to competitive equilibria, it’s worth noting the negative result as well: if objects are not substitues to everyone, market-clearing prices may fail to exist. A simple two-good, two-bidder example:

Bidder	v_A	v_B	v_{A+B}
1	a	b	$a+b+c$
2	$a+0.6c$	$b+0.6c$	$a+b$

Competitive equilibria are efficient, so bidder 1 must get both items. For this to happen competitively, we need both $p_A \geq a+0.6c$ and $p_B \geq b+0.6c$, so that bidder 2 does not demand either object. But then $p_A + p_B \geq a+b+1.2c > v^1_{A+B}$, so bidder 1 would not demand anything.

Milgrom gives a general proof that if there are at least three bidders, and the set of possible individual valuations includes all substitutes valuations and at least one other, then there will be realizations such that no competitive equilibrium exists. (There are other analogous results as well.)

Empirically, there have been simultaneous ascending auctions that have been regarded as very successful despite perceived complementarities across items – in early U.S. and Mexican spectrum auctions, regional players managed to assemble the packages of complementary licenses they needed. In settings where all bidders regard two items as complements, the problem is generally not severe – whoever wins one of them, will likely want to win the other, so the exposure problem is not severe. The problem arises when items are complements to one (or more) bidders, but substitutes to others.

On the other hand, the results of a 1998 spectrum auction in the Netherlands highlight the magnitude of the exposure problem. Eighteen lots were offered for sale. Two were designed to be large enough for a new entrant to establish a new wireless phone business. The other sixteen were too small to be valuable alone to new entrants, but could be used by incumbents to expand their systems; or a new entrant could assemble a collection of several licenses to launch a new business. So the smaller licenses were likely to be complements for new entrants, but substitutes for incumbents.

The final price per unit of bandwidth ended up being more than twice as high for the two large lots as for any of the smaller lots. Even though they might have been able to enter more cheaply by bidding on multiple small lots, the entrants may have avoided the exposure problem by simply bidding more aggressively on the large lots. (Although a footnote in Milgrom suggests other differences in the licenses may explain some of the price difference.)

Clock Auctions

A slight variation on a simultaneous ascending auction is a clock auction. The main motivation for using a clock auction is that it takes less time - in a simultaneous ascending auction, even after "most" of the allocation is determined, bidding continues on a few remaining items for many additional rounds; but since bidding is still open on all items, the other winners can't take their prizes and walk away. Clock auctions give a way to "speed up" convergence to the competitive outcome.

The clock auction works like this. Prices begin at some low level, and bidders submit their demand at the level. (Assume for the moment that at all relevant prices, demand is single-valued.) The auctioneer figures out which items are **overdemanded**, that is, which items are demanded by more than one bidder, and raises the prices of those by some discrete amount. Bidders submit their new demands, and so on. The process resembles that of Walrasian "tatonnement" – adjusting prices until markets clear.

There are, of course, some logistical problems. First of all, when prices are raised a discrete amount, they may overshoot – an item that was overdemanded may become underdemanded. One procedure for dealing with this is used in French electricity markets. The auctioneer announces his intention to raise prices on some set of items (in this case, electricity contracts) by some discrete amount δ. Given that information, the bidders give instructions to electronic bidding agents, telling that at what prices to reduce their demand for particular contracts. Once these instructions are given, the prices are raised continuously, and only the electronic agents bid; thus, prices can stop going up when particular contracts stop being overdemanded.

(Clock auctions can be run for distinct items, or for identical items, such as standard electricity forward contracts or financial instruments, where demand is for a quantity, not particular items.)

Clock auctions, however, turn out not to be immune to the same types of low-revenue equilibria we saw before in uniform-price auctions.

Vickrey Auction

Next, we consider a generalization of the second-price auction, the Vickrey auction, or Vickrey-Clarke-Groves Mechanism, or the Pivot Mechanism. The Vickrey auction is basically designed to be a direct-revelation mechanism for a setting with multiple items, or multiple units of the same good: bidders report their full set of preferences (how much they value each possible combination of goods), and the auctioneer allocates the objects efficiently, and charges the prices that make truthful revelation a dominant strategy.

Suppose there are L (distinct) goods for sale to N bidders. It's customary to think of the seller as bidder 0. Let x be an allocation of the goods to the bidders. (Think of x as an $L \times (N+1)$ matrix of 1's and 0's, with the constraint that each row sums to 1.) Then x^j is the set of goods allocated to bidder j (or kept by the seller, when $j = 0$).

Bidder preferences (and the seller's valuation for the unsold objects) are

$$v^j : 2^{|L|} \to \Re^+$$

Given that each bidder j reports preferences $\hat{v}^j$, the Vickrey auction gives the efficient allocation,

$$\hat{x} = \max_x \sum_{j=0}^{N} \hat{v}^j(x^j)$$

The payment each bidder makes can be thought of as the externality his presence imposes on everyone else; that is,

$$p_i = \sum_{j \neq i} \hat{v}^j(\hat{x}^j) - \max_x \sum_{j \neq i} \hat{v}^j(x^j)$$

The first term is the value (before transfers) realized by the other players under the chosen allocation $\hat{x}$; the second is the value they would realize if bidder i had stayed home.

It's not hard to show that reporting truthfully is a weakly dominant strategy. The Vickrey auction is guaranteed to be efficient, and can be shown to be the **only** efficient mechanism. However, the Vickrey auction also has lots of significant problems.

- First of all, the Vickrey auction requires each bidder to calculate, and communicate, his valuation for every possible package. So if there are L objects for sale, the bidders must each evaluate $2^L - 1$ different packages. For L large, this is obviously unlikely.

- (Larry Ausubel, in "An Efficient Ascending-Bid Auction for Multiple Objects," gives a clever design for an ascending auction for many identical objects which avoids this problem; it implements the Vickrey outcome when values are private, and may outperform it with interdependent values, but only works with homogeneous objects.)

- Second, bidders no longer have a dominant strategy if they have a limited budget. The Milgrom book gives a simple example where a budget-constrained bidder's best-response depends on his opponent's reported preferences.

- Third, in many settings, the Vickrey auction leads to low revenue for the seller. (We'll make this more precise in a minute.)

- Fourth, and most interestingly, if the goods are not all substitutes for all bidders, the Vickrey auction exhibits certain types of very flukey behavior:

  - Revenue can decrease as more bidders are added
  - Although (by definition) no single player has a profitable unilateral deviation, it's possible for a group of **losing** bidders (bidders who get nothing in equilibrium) to jointly raise their bids and end up winning objects at profitable prices
  - Bidders can benefit from shill bidding (hiring a friend to bid "against" them)

A simple example that will illustrate several of these problems. There are two objects for sale, the seller values them at 0, and four bidders, two of whom see the objects as complements, two of who do not:

Bidder	v_A	v_B	v_{A+B}
1	0	0	1000
2	0	0	900
3	1000	0	1000
4	0	1000	1000

The Vickrey auction awards object A to bidder 3, and object B to bidder 4, and both pay nothing! To see this, calculate bidder 3's payment. Given the efficient allocation, his three opponents get a total value of 1000; if he wasn't there, his three opponents would still get a total value of 1000. So the Vickrey auction can give 0 revenue even when there seems to be plenty of competition for the objects. (When objects are substitutes, the Vickrey auction can still lead to low revenue.)

The same example shows how revenue can decrease in the number of bidders. If the auction were held just between bidders 1 and 2, bidder 1 would win, and pay 900. Adding bidders 3 and 4 drops revenue to 0.

To see loser collusion, modify the example:

Bidder	v_A	v_B	v_{A+B}
1	0	0	1000
2	0	0	900
3	400	0	400
4	0	400	400

Now, the dominant strategy equilibrium allocates both objects to bidder 1; bidders 3 and 4 get nothing. But if they both raised their bids to 1000, as we saw before, they'd each get an object, and pay nothing.

Finally, to see shill bidding, modify the same example once more:

Bidder	v_A	v_B	v_{A+B}
1	0	0	1000
2	0	0	900
3	0	0	800

Clearly, bidder 3 gets nothing. But if he reported his value as $1000, 0, 1000$, and got a friend to enter and report $0, 1000, 1000$, together they'd win both objects and pay nothing.

The Milgrom book contains a theorem saying, basically, that all these problems go away if the objects are substitutes.

(He also shows that when valuations are substitutes, the Vickrey outcome is in the core of the game. The core of a game is basically the set of outcomes that are coalition-proof. That is, an outcome of a game is in the core if there is no coalition of players who can deviate from the outcome and get higher payoffs by transacting by themselves. Obviously, any deviating coalition in this setting would have to include the seller.)

So the upshot: the Vickrey auction works well when goods are substitutes, poorly when goods are not substitutes. ("Poorly" means, revenue can be very low, and the auction can be manipulated.)

First-Price Package Auction

Finally, there is the analog of the first-price auction in a multi-good setting: the first-price package auction. Bidders submit sealed bids for whichever collections of goods they're interested in; the seller accepts the collection of bids that maximizes his revenue, and bidders pay their bids for the package they receive. (These are also referred to as "menu auctions", since each bidder offers a menu of bids to the seller.)

The London bus route auction has used this design, as have some others.

Unfortunately, package auctions have so far proven too complex to analyze as Bayesian games. That is, no one has gotten all that far developing general results on package auctions with private information. Bernheim and Whinston studied the full-information case, that is, the case where all bidders' valuation functions are common knowledge. (This is, of course, a pretty frustrating limitation of the model.) They study a particular type of equilibrium of this game, and show that the outcomes mostly match up with the set of bidder-optimal core outcomes.

Threshold problem (example).

Ausubel-Milgrom Ascending Proxy Auction

Larry Ausubel and Paul Milgrom introduced an ascending-auction version of a package auction, using proxy bidders, that is, computer programs that bid for you after you tell it what to do.

The auction works like this. Each bidder reports a valuation function $\tilde{v}^j : 2^L \to \Re^+$ to the his electronic proxy. The auction takes place in a series of rounds. In each round, the seller chooses the feasible allocation x that solves the problem

$$\max_x \left(\sum_{j \in N} \beta^j(x^j) + v^0(x^0) \right)$$

where $\beta^j(x^j)$ is the highest bid that bidder j has submitted (at any point) for package x^j. Bidders who are allocated some package are called provisionally winning bidders.

In each round, provisionally winning bidders do not submit new bids. The other bidders can submit a single new bid: either δ for a package they haven't yet bid on, or $\beta^j(x) + \delta$ for a package x they'd already bid β^j for, where δ is the bid increment, which is assumed to be small. Call these minimum bids $m^j(x)$.

The proxy for the bidders who are not provisional winners choose the package x that gives the highest profit at the minimum bid, that is, the package that solves

$$\max_x \tilde{v}^j(x) - m^j(x)$$

and submits the minimum allowed new bid on that package (assuming this potential profit is positive). Rinse, repeat.

Ausubel and Milgrom show that if goods are substitues, everyone submitting their true preferences is an equilibrium, and the resulting payoffs are the unique bidder-optimal core outcome. There are, however, other equilibria as well.

We've been working in a full-information environment, but with substitutes, truthful reporting is a best-response to any opponent preference profile, and therefore to any opponent strategy, which means that truthful reporting would remain a dominant strategy if we added private information. But without substitutes, this would break down.

Conclusion

So that's a very brief look at a bunch of formats for multi-unit auctions. Pretty consistently, we find that multi-unit auctions tend to behave well when goods are substitutes, but face a host of problems when goods are complements.

First-price auctions with package bidding are generally too complicated to solve for equilibrium. There are a few exceptions – people who have taken very specific examples and solved them. (I have a paper like this with a coauthor, showing a very simple multi-unit setting where we get non-monotonic symmetric equilibria.) There's also been some interesting experimental work done – to see how people actually behave in these types of auctions.

Given the complexity – Vickrey auctions require reporting your valuation for 2^L different bundles, and first-price auctions require bidders to solve for an equilibrium that game theorists haven't been able to solve – I'd expect actual outcomes in these auctions to be very sensitive to the precise details – the interface, how bidders are prompted to act, and so on. Hence, the need for experimental work. And of course there's lots of room for theoretical advances here, they're just, you know, hard.

0.1 If there's time...

One more note. These were all auctions for settings in which bidders want more than one object. Even when bidders have unit demand, if there are multiple goods, things can still be tricky. There's a working paper by Milgrom and Weber, "A Theory of Auctions and Competitive Bidding, II," written in 1982, meant as a companion paper to the other auction stuff they were doing at the time; however, they got stuck on some of the proofs, which remain conjectures. As a result, the paper was passed around, but not published, until 2000 when it was included in a book pulled together by Klemperer.

The general setup: n bidders, k objects for sale, but instead of all being sold in one auction, consider the possibility of a series of single-good auctions. The objects are assumed to be identical, so bidders draw a single type at the beginning and don't care which auction they win (subject to what they pay and the information they get along the way). Each bidder only wants 1 item. And valuations are the usual Milgrom/Weber affiliated/interdependent.

To very briefly summarize their findings. (These are proven for some of the auction formats, conjectured for the others.)

- With IPV, sequential auctions are revenue-equivalent to a big one-shot auction; with sequential auctions, winning bids (price paid in each auction) are a Martingale, that is, the expected value of the next one is the same price as the last one
- With affiliation and interdependent values, sequential auctions revenue-dominate a one-shot auction, and winning bids on average should have positive drift – on average, they should increase from auction to auction

So we have some results which are proven for sequential first-price auctions with price announcements, and posited to hold for sequential auctions more generally:

- Sequential auctions should revenue-dominate the corresponding one-shot k-unit auction
- Prices paid should drift upwards over the course of the sequence

The first may be tough to verify, but the second should not. This is the other thing that went wrong with this paper: the cleanest empirical observations they could find directly contradicted it.

In 1981, RCA planned to launch a satellite carrying 24 transponders, each able to relay one cable TV signal. A bunch of cable companies were interested in leasing rights to these transponders. RCA decided that seven of them would be sold via sequential auctions, run by Sotheby's. The auctions were English auctions, with the winning price visible after each auction. (Total revenue was over $90 million.)

The situation fit the model pretty well: the items for sale were clearly identical, and it's reasonable to think that each cable company would only want one. So we would hope that the theoretical results would fit the data.

Here are the prices paid in each of the seven auctions, in millions: \$14.4, 14.1, 13.7, 13.5, 12.5, 10.7, 11.2. To the naked eye, these do not appear to be realizations of a random process with an upward drift.

(Milgrom and Weber do give a story for why this happened. It's basically a principal-agent problem. Suppose each firm appointed an agent to bid for them, and gave them the instructions: "If you can get a lease for less than X, do it; otherwise, don't." In that case, the agent could "fail" in two ways: paying more than X, or not getting an object when one of them sold for less than X. The only way to avoid any risk of failure would be to bid X every time. (Thus, the agent actually making the bid was not maximizing expected profits, but minimizing the risk of looking bad to their bosses.) If everyone did this, then the firm with the highest remaining "limit" would win each round, and the price would go down each time.

MW point out that, if this were the case, the first few firms to win would "overpay", and would therefore regret the sale and, given the choice, might choose to undo it. As it happens, this opportunity arose: the FCC threw out the results of this auction, saying that since identical leases sold at different prices, this constituted discriminatory pricing. RCA then offered the leases for sale on a first-come, first-served basis for a price of \$13 million each; several of the firms that had bid above \$13 million in the auction declined to pay that price.)