[SQUEAKING]

[RUSTLING]

[CLICKING]

GLENN ELLISON: OK, so let me get started. Today I'm going to be talking about auctions. And let me say that this lecture is going to be-- there's a lot of math on the slides. I like to have lectures where you watch everything and you understand it. Today there are going to be some of those slides that you just look at them at home afterwards to try to get how I'm doing some of the derivations.

But anyway, we're used to and I think almost all this class, we've been of talking about firms and the firms have been setting prices and the firms have been setting prices to maximize profits or some kind of oligopoly competition. There are many real world applications where firms do not set prices. And instead, firms use auctions as a mechanism to set prices.

Three of the biggest applications. Government procurement. Whenever the government is hiring contractors to pave roads or fix bridges or things like that, these are usually done through procurement auctions where the government asks firms to submit bids for how much they would need to charge to pave a road and then choose the winning bidder to be the lowest bidder.

In financial markets, the US auctions Treasury bills. And Treasury bills are trillions of dollars, and so there are trillions of dollars in auctions and that thing. Online advertising is also absolutely enormous. Every time you're looking at a computer screen and you see some ads on the page or whatever, whether it's a news site or a search engine or whatever, almost all of those ads these days are being auctioned off rather than being sold, because it's just very hard to know what the price is. What is it worth to show you an ad at a particular point in time? And so there's auction mechanisms that are setting the prices.

There have been lots of academic research in the past few decades. I think one of the reasons why people in IO like this is because we always like to write down models and have these mathematical models that shed clarity on something. And often in many other applications, we're making up the model and making up the game that people are playing. In auctions, it's very clear that there is a real game and you can see the game that people are playing. It helps you do this nice connection between theory and empirics.

And in terms of what's happened in IO, this really grew in the-- I guess that was in the 1990s or early 2000s when formally the airwaves that are used for cell phone signals and for TV licenses and things like that, in the old days, those were simply given away by the government. And if the government gave you one of the three licenses to conduct television in a city, you would make enormous amounts of money.

At some point, the government decided to shift and auction off the right to provide cellular telephone service or auction off the right to send TV signals over the air. Those auctions raised tens of billions of dollars. Therefore, people wanted to hire lots of auction theorists to work on the auctions and design the auctions for them. And they raised a lot of interest.

More recently, the online advertising example, again, there's tremendously complex auctions that you can run to auction advertising, and it spurred a lot of interest. There's also been a big literature on structural estimation of auctions. And mostly Tobias is going to do that in Wednesday's lecture.

So I'm going to start with what's not a great model for applications, but it's a great model for solving the model and getting insights from it. So it's a special case, but I will do a lot of my analysis in the independent private values model. So here we have bidders indexed by i equals 1 to n. There's one object to be sold. So you can think of this as people, they're at a charity fundraiser and they're auctioning off a bottle of wine. And whoever wins the auction is going to take the bottle of wine and drink it or something like that.

Each bidder observes a signal si that relates to their value. We're going to assume that these signals are independent across bidders. The term private values means that I know what the good is worth to me. My value is a function only of my own signal, not of everybody else's signal. So I know I'm an expert in wine. I see the wine. I see the label. I just know how much I'm going to enjoy drinking that bottle of wine. So my value depends only on my own signal.

Once you make the decision that-- once you have the private values model, where your value only depends on your own signal, it's without loss of generality to say the signal is your value. So that's what we're going to typically do is everyone just observes their own value. We're going to assume that bidders are risk neutral. That is, you get a value of s minus p if you win the auction at a price of p. If you were to lose and be forced to pay p, either as an entry fee or something else, then your utility would be negative p. And the bidders are only going to observe their signal.

This is also a model that people use for procurement auctions. If you think about it, you have a government is inviting contractors to bid on repaving Memorial Drive. You can think of the s as the amount of money that it will cost the firm to do the paving job. And s is obviously a big negative number. If you're asked to pave a road, it will cost you $2.7 million. So the s is therefore negative $2.7 million. And then what you're bidding is you're bidding a price. The price is going to be negative, like the price could be negative $3 million. And so if the price of getting the contract is negative $3 million, your cost is negative $2.7 million, you get $300,000 in profits.

First auction that's considered always by theorists, not so often used in practice, is the second price sealed bid auction. In the second price sealed bid auction, what you would do is you'd have every contractor say, I will pave the road for $2.5 million. I will pave the road for $2.7. I would do it for $2.8. Everyone just puts their bid in an envelope, seals it. They're all handed in. And then on the day of the auction, all the envelopes are opened. Whoever has bid, well, in the bidding for a positive price, whoever bids the highest amount is going to win. But instead of paying what they bid, they pay the second highest bid.

So if the we're bidding for a painting and the bids are $2.4 million, $2.3 million, $2.2 million, you open the envelopes. This person gets the good, but they only pay the second highest bid instead of paying what they bid. Why do we do that? It turns out it becomes a very easy auction to analyze. And so it's a very nice mechanism in that sense.

Proposition: in the second price sealed bid auction, it's a weakly dominant strategy to bid your true value. So you don't have to think strategically about what other people are bidding, what's the good worth to them. You just have this dominant strategy, which is just to write down your own value, bid that, and that's a dominant strategy equilibrium.

A good way to think about that is to just think about what are the deviations you could do? So suppose your value is si. You consider writing down something that's different from your value and let b hat be the largest value written down by everybody else. There's six possible orderings of bs and b hat.

In the middle cases here, I have the cases where-- sorry, in the first top and bottom cases. In the top and bottom cases, when somebody else has written down a number that's bigger than both your value and the number you wrote down, well then you lose whether you wrote down si or wrote down bi, and so your value is 0 either way. The other case is if everyone else bid lower numbers, things that are lower both than your value and the thing you wrote down, in this case, you're going to win the auction whether you wrote down bi or si and you're going to pay b hat either way. So the fact that you've done something different from writing down si also doesn't make a difference.

The two cases we're writing something different from si does make a difference is here and here. Suppose that your value is si. Someone else wrote down a number that's bigger than your value and you write down something even bigger. Well then if you've written down your true value, you would have lost. Now that you win, you get the object and you pay b hat, which is more than it's worth to you. So you end up regretting that you won, because you've bid something more than it's worth.

And then the final case is I would have won if I'd bid si, but I write something lower, and then I lose. And so here, instead of getting utility si minus b hat, I get utility 0. So in all six cases, I'm weakly worse off and in two of them I'm strictly worse off. So just writing down my true value is an optimal strategy.

It's not the unique equilibrium. For instance, you could have one person bid in this model. One person bids s upper bar and everybody else bids 0. And then all the bidders bidding 0 are happy to bid 0, because they can't win unless they bid s upper bar, which is more than it's worth to them. The person bidding s upper bar is happy to bid s upper bar, because he's going to get it at a price of 0.

Usually we don't pay attention to those equilibria. So usually we think about this equilibrium of this game where everyone just plays their dominant strategy. And that's what they do. It is the unique symmetric equilibrium if you have continuous distributions.

So one thing that's nice about this dominant strategy is now it's easy to figure out how much revenue the seller makes from this auction. So the amount the seller makes, I'm going to write sn minus 1 colon n for the n minus first order statistic of n draws from this distribution. So this is the n minus first highest. So because whatever the bids are, you always get the second highest or n minus first lowest bid. We just need to compute the expected value of this second highest draw from the distribution.

And a good way to think about that is, how is it going to be that the second highest bid is exactly $2.3 million? The way that's going to happen is, suppose we have one bidder like bidder number three who has a higher value than this, one bidder who bids exactly this or within dx of this, and then you have all the rest of the bidders down here. There are many ways this can happen.

It doesn't have to be bidder three up here and bidder seven here. It could be bidder four here and bidder nine here. It could be bidder one here, bidder two here. But for this particular configuration, the probability that bidder three is above this point x is 1 minus capital f of x. The probability that bidder seven is in this tiny interval is f of x dx. Because it has to be an interval dx around this. And then the probability that all the rest of the people are here is f of x to the n minus 2.

So the total probability of a configuration like this with the second highest bid is x is just n times n minus 1 times this times this times that. Because there are n people who could be the high bidder, n minus 1 who could then be the second highest bidder. So this thing here and n minus 1 capital f of x to the n minus 2 little f of x1 minus f of x dx, that's the density of the second highest value. And so the expected second highest value is just integrate x times that density.

So just as an example here, I do this for the uniform distribution. And the uniform distribution, this says that an auction with n bidders raises n minus 1 over n plus 1 as revenue. So the revenue from a second price sealed bid auction goes to 1 as the number of bidders goes to infinity, because the second highest value is almost the first value, which is almost the top anyway.

And this density calculation, it's an example of just how you compute a k-th order statistic, which is you have n minus 1 choose k minus 1 possible groups of people who can be above x. Who can be, sorry, below x. f of x to the k minus 1 is then the probability that each of them has a value below x. 1 minus f of x to the n minus k is a probability that everybody else is above x. And then you have f of x dx is the probability that someone has a value exactly x.

And so computing expected values of auctions. Revenues, you're often computing these order statistics. And when the draws are independent and people bid truthfully, that's a fairly easy calculation to do.

A second auction that's-- second price sealed bid auctions are this nice economic gain, nice model, because they're very simple. In practice, we don't see auctions conducted that way. In practice, what you often see is auctions conducted with an oral ascending bid process where you-- anyway, you probably haven't been there, but you watch people do auctions on TV for paintings and it's always people with paddles up in the air who will bid $1 million, $1.1, do I hear $1.2? People just keep raising their hand every time they want to bid higher. And then they keep bidding higher until the guy says, going, going, gone, because no one else is willing to bid more than the current bid amount.

Analyzing that kind of game where people keep raising their hand would be a very complicated, extensive form. So what people do instead is simplify to this what's called the button auction design, which is the clock is going up and the auctioneer keeps calling out higher and higher numbers. And you just hold your finger on the button as long as you still want to buy the good at that price.

And then once you take your finger off the button, you're out of the bidding and you've dropped out forever. So it's like, who's in, who's in, who's in. And then you just take your finger off and you're out permanently. That's nice, because it converts it to a basic one dimensional strategy space, which is I just choose at what point do I drop out of the auction?

Thinking about that game of at what point or if I was giving instructions to a computer, at what point should you drop out of the bidding? It's clear that in some sense, this game is the second price sealed bid auction, just in a different form. So a perfect Bayesian equilibrium of this button auction model is just everybody drops out when the auctioneer gets to their true value. So as long as he's saying if the good is worth $2.4 million to you, as long as he's saying $2 million, $2.1, $2.3, you keep your finger on the button. And as soon as he gets $2.4, you take your finger off the button and say, I don't want it anymore.

Why is this true? I think this one is actually it's even more obvious than the second price sealed bid auction. What's my dominant strategy? As long as I want to win, I keep saying I'm in. As soon as I don't want to pay what's being charged, I drop out. So it's the same argument as a second price sealed bid auction. Dropping out earlier could only make me lose something that I would want to have bought. And bidding past my value, all I can do is lose and end up paying more than I wanted to get for it. So the game is still the same. We get the second highest price.

In practice, what's funny is if you run second price sealed bid auction models, people often get it wrong. People understand that I won't actually-- people understand I won't actually need to pay what I bid. I'll probably have to pay something less. They don't do that full consideration of the six cases and say, OK, I should still bid exactly what it's worth. And you will see in second price experiments, people over bid in second price sealed bid auctions. They realize, OK, I'm probably not going to have to pay what I bid, so therefore I can bid more than it's actually worth to me, and I probably won't have to pay that. Somehow the conditioning logic doesn't make it through.

But in practice, I would say in practice people will tend to do first price-- will tend to do auctions of this variety rather than the second price sealed bid auction. One thing that is different about the auction from the auctioneer's perspective is the information that's revealed about the bids. In this auction, when you open the envelopes, you learn what the good was worth to everyone, and you may learn how much surplus went to the leading bidder, because the second highest bidder had a value much below the highest bidder. Whereas in this auction, you never learn. You learn when everybody else drops out. You never learn the value of the highest bidder.

There were some cases. There was a famous spectrum auction case where they were auctioning off I think it was cellular telephone service, right to conduct cellular telephone service in some small town in New Zealand that basically only one person wanted to provide the service there. And I'll make up these numbers. They're completely wrong.

But something like the high first person submitted $37 million and the second one submitted $15,000 and then everyone else submitted 0. And then they open up the envelopes and they're like, OK, your brilliant auction design says we give it to the bidder number one, who is willing to pay $37 million. We charge them $15,000 and everyone's like, that was the worst auction design in history. How could you have recommended that?

Obviously, if we'd done the oral ascending auction, all we would have seen is that only two people show up and one of them puts his hand down at 15,000. And we're like, OK, it's not worth providing service in this town. That was a reasonable price that we got. And so in practice that ex-post embarrassment of the auctioneer may be a factor that leads us not to do the second price sealed bid mechanism.

Another mechanism that we see even more often in practice is the first price sealed bid mechanism. So the first price sealed bid mechanism would be like the second price sealed bid. You're running your procurement auction. You ask all the bidders to bid what price you would need to give them to get them to pave the road. They all submit their bids. And then the highest bidder wins and actually pays the bid. Or if we're buying something, I write down I'm willing to pay $2.4 million for the painting. Someone else writes down $2.3, $2.2. We open the envelopes. If I wrote down, I'm willing to pay $2.4 million, you charge me $2.4 million.

What's the equilibrium of this game? Here this is a game where you have to shade your bids. You can't bid what the good is truly worth to you, because sometimes bidding what the good is worth to you is a weakly dominated strategy. If I'm willing to pay exactly $2.4 million, at any higher price, I wouldn't want it. If I bid $2.4 million, then if I lose, I get 0 utility. And if I win, I get 0 utility, because I've bid exactly everything that the good was worth to me. I get 0 surplus.

So what do you do? It's like monopoly pricing where you shade your bid and you decide, I'm going to bid less than my true value. As I go further and further down below my true value, I get more profit if I am chosen as the winner, but it's less and less likely I'm going to win. So it's kind of like the monopolist raising price above cost. I get more and more profit if I do make the sale, but it's less and less likely that I make the sale. And you make that trade off of probability of winning times surplus versus surplus if I win.

So let's think about what the equilibrium would look like. So suppose we had a symmetric equilibrium. Everyone's using these strictly increasing strategies. b star of s. The higher is my value or signal s, the higher I bid. My payoff. If I bid bi and everyone else uses b star is I get s minus bi is my prize. I get that prize if everybody else bids less than b star of sj is less than bi for all j. So I win if everybody else bids less than the amount that I bid.

So that tells me that maximization problem would be that my optimal bid has to be the maximizer of si minus bi times f to the n minus 1. So the probability that b star of sj is less than or equal to bi is just the probability that sj is less than or equal to b star inverse of bi. So the probability that everybody else's signal is below b star inverse of bi is just f of b star inverse of bi raised to the n minus 1 power.

And a trick for simplifying the analysis of this that you would think shouldn't work, but this is kind of like what we do in mechanism design often, is instead of thinking about people choosing their bid, think of people choosing pretending to be some other type and bidding like some other person with some other value bids in this.

So if I'm thinking about what type s do I want to imitate? I'm then going to imitate type s by choosing b star of s. And then my value is going to be my price. I'm going to win s. I'm going to get a payoff of s minus b star of s if everyone else's value is less than s. So if I bid like the bidder who has type s, then I win whenever everybody else's type is less than s. And so I've eliminated this inverse function here, and I've got my payoff function I'm trying to maximize is s minus b star of s times f to the n minus 1 of s.

So on the bottom line of that screen, I've taken the first order condition differentiating this with respect to s. Differentiating this with respect to s and setting it equal to 0 at s equals si. And so the derivative of this with respect to s is this term times the derivative of this one, which is n minus 1 f to the n minus 2 of si times little f of si and then plus this times the derivative of this, which is just minus b star prime evaluated at si. So that's the first order condition for my maximization when I'm thinking about which type do I want to copy and noting that the type I copy has to be me.

So the trick here is that I know the equilibrium. I know the equilibrium is s star equals si. And so because you know the equilibrium is the identity, that simplifies what the expression is for the first order condition evaluated at the equilibrium.

And now in the next three slides, I'm going to do three different rewritings of this. One rewriting I'm going to do is just. Think of myself as solving this for the bid. So I'm going to solve this equation for b star, and I'm going to do that by just taking this expression b star and n minus 1 f to the n minus 2 little f, move it to the other side, and then divide by this and see what's left. And what you get is the optimal bid is your true value minus a markdown. So this is an expression for how much you shade your bid.

If I look at the numerator, you'll see that the numerator of this f to the n minus 1 of si, that's just the probability that everyone else bids less than I bid. Because that's the probability that everyone else has a lower type than I do. So this is the probability that everyone else bids less than what I'm bidding in equilibrium. And then the denominator is just the derivative of the numerator with a change of variables formula. So you differentiate this g with respect to this. This is what you get. And then you have 1 over b star prime from a change of variables formula.

So what we get is that what you do in a second price sealed bid auction is you mark down your bid. You mark it down by the ratio of the CDF of the bid distribution, the highest bid distribution evaluated at si, divided by the PDF of the highest bid distribution.

So when do you shade your bids more? You're more aggressive in shading your bids if the probability of winning is higher and the density of the rival bid distribution is lower. Because this is kind of like an inverse elasticity rule for markups. The more bidders there who are bidding around almost exactly b star of si, the more you don't want to shade your bid, because if you shade your bid, you're likely to lose.

But if this density is low, then every extra bit of shading you do has very little-- there's very little demand elasticity. There's very little impact on your probability of winning. So then you shade more aggressively. And so bid shading is based on how likely you are to win, which gives you the inframarginal gains from lowering the bid. And then the g is what's the increased probability of losing. And so I mark down based on those ratios.

And this expression, I did it because Tobias is going to be lecturing on Wednesday. This is also a very, very big part of the structural estimation of auction models, because this is just the CDF of the rival bid distribution. This is the PDF of the rival bid distribution. These are estimable objects. That basic insight of structural estimation of auctions is this is observable, this is observable, this is observable. So that's how you can back out people's values. Anyway, Tobias will do that in much more detail on Wednesday.

Second thing I'm going to do with that first order condition is I'm actually going to solve this auction. And again, this is something that it should not work, but the model is elegantly designed so that it does work. So this was our first order condition that determined the optimal bid. I'm going to rearrange it by putting all the things with negative signs on the left side.

So I put the b star prime term on the left side, and I put this b star times this stuff on the left side. So I get b star of si times this order statistic plus b star prime times that CDF equals si times this density. And then you look at that expression and you're like, wow, I can integrate both sides of that. The left side is an exact derivative. It's a derivative of this expression, b star times f to the n minus 1, because this is b star times the derivative of f to the n minus 1 and this is f of n minus 1 times b star prime. So b star times f to the n minus 1 is the integral from s lower bar to si of this density on the right side.

And then again, skipping steps, you can integrate this expression on the right side by parts. That's a trick someone found a while ago. And you can get a closed form for expression, which is your bid is your markup, your bid is your value minus this markdown, which is the integral from s bar to si f to the n minus 1 of x ex over capital f to the n minus 1 of s.

So this first price sealed bid auction is solvable. The only thing you need to plug in is what is the distribution of the values. Is it a uniform distribution? Is it an exponential distribution? Any distribution for which you can do this integral of f to the n minus 1, you can just get a nice closed form solution for what an auction will be.

For instance, down here, I did it for the uniform distribution. In the uniform distribution, f of x is just x. And so it's just the integral x to the n minus 1 dx. That integrates to x to the n over n. You evaluate that and this tells me that the optimal bid with uniform distributions is n minus 1 over n times your value. So as the number of bidders gets large, you're always shading your value.

And with the uniform distribution, you're shading your value in a linear way. With n bidders, you just bid n minus 1 over n times your value. So if n is two, you bid half your value. n is 3, you bid 2/3 of your value. n is 10, you bid 90% of your value. But it's nothing special about the uniform distribution. You can solve this model however you do this.

Another thing that you get from this is a proposition that helps you understand why the auction raised the revenue it does. The optimal bid expression I wrote down on the previous page is just I bid the expected second highest value conditional on having the highest value. So I know my signal. I do the thought experiment. If I did have the highest value, so everyone else's value was less than mine, what would be the second highest value? I bid the expectation of that.

Why is that true? I don't have great intuition, but it's easy to see it mathematically. So in the previous slide, I had this expression for the-- I had this thing. I recognized both sides were derivatives, and I got this expression for b star of s as an integral. So that gives me b star of s is this integral of this density on the top and f to the n minus 1 on the bottom.

But if you look at what this density is, n minus 1 f to the n minus 2 little f of x dx, that's the density of the highest of n minus 1 draws from f. And so I'm integrating x times the density of the highest of n minus 1 draws from x from s lower bar up to capital s, and then I'm dividing by f to the n minus 1 of x. So when I divide by f to the n minus 1 of x, this is giving me this density conditional on the highest value being less than si. And so this is just the expected second highest bid conditional on the second highest bid being less than si.

So third way to think about bidding. So thought about-- thinking about bidding in terms of markups. It turns out that a way to think about what the markups imply is that the way you bid is you just take your value, think about what the second highest value would be, and you bid that in expectation. So corollary of that observation is that first and second price auctions yield exactly the same expected revenue to the seller.

The second price auction, clearly you're always getting the second highest value as your revenue. So it's just the expected second highest value is the revenue. In the first price sealed bid auction, what you're getting is the expectation over the-- you always sell it to the high value person. So I'm getting the expectation over the value of the high value person of the expected second, the expected highest of-- the expected second highest value conditional on the second highest value being less than that level.

And so iterated expectation says that's just another way to compute the expected second highest value. I just condition on the highest value and say, what's the expected second highest value conditional on that being the highest value? And then I integrate out over all possible highest values. That is the expected second highest value. So because everyone's bidding expected second highest values conditional on winning, you get the same expected second highest value as the revenue.

So this is a special case of what's called the revenue equivalence theorem that you have these seemingly different auction mechanisms, one where people bid truthfully, one where people shade bids, and somehow they raise exactly the same revenue. I would say this is actually one of the most misunderstood things.

You talk to people who are doing selling auctions online, which are often sold through second price mechanisms. And just at every company that does this, many times per year, every year, some executive comes up and realizes, wait, people are bidding this and we're only charging them the second highest value. I can get this 1 over n thing. There's this big gap that's $42 million. We should switch to a first price auction instead of a second price auction.

And obviously, the mistake in that reasoning is yes, if you had the same bids and you changed the rule, you would get more money. But when you change the rules, people will just start shading their bids. They will bid less. You'll raise exactly the same revenue with the first or second price mechanism. We often think second price mechanisms are better because in second rights mechanisms, people can think about how much to bid, which is just what's it worth to me, whereas in the first price mechanism, they have to do this complicated reasoning about bid shading and what would be the expected second highest bid given what I know. It's putting all that work on the firms is not doing you any good.

Anyway, revenue equivalence theorem-- and I'm going to go through this quickly. Consider more general auction mechanism. Imagine you had any game. The bidders just submit bids b1 through bn. The good is allocated to bidders with some probability that depends on what all the bids are. And the payments are some function of what all the bids are. So the first price sealed bid auction would fit this. You give it with probability 1, the person with the highest bid, and they paid their bid. The second price auction is, again, you give it to the person who had the highest bid, but they pay the second highest of these numbers. But it's a general class of auction mechanisms.

Theorem is consider the symmetric independent private values model. Suppose that a general auction mechanism has an equilibrium bidding function in which the object is awarded to the bidder with the highest value and a bidder with value s lower bar gets 0 surplus.

Then in equilibrium, the seller's expected revenue is the same thing we've just seen, the expected second highest value. And not only are the auctions equivalent in terms of revenue, they're equivalent from the bidder perspective in that the expected utility of a bidder who has type si is just the integral from s lower bar to si f to the n minus 1 of x dx. You'll hear people refer to this as the information rent of the bidder.

But in any mechanism you choose that has these two natural properties, you always raise the same revenue, and you raise the same revenue because you always maximize gross surplus and you give away this much surplus to a bidder of type si. So they're strongly equivalent in that they raise the same revenue for the seller and they give the same expected utility for a buyer and expected utility, even conditional on knowing your type, it's always the same.

And the thing I gave earlier on first and second price auctions is a special case of this. Because a first and second price auction games, I gave you the option, the object was always being awarded in equilibrium to the bidder with the highest value, whether in the second price auction because you stated your value or in the first price auction because you bid n minus 1 over n times your value. And then also a bidder with the lowest value was going to get 0 profits, because they win with probability 0. Yes.

STUDENT: Isn't the second condition, more of a result? Because I thought the whole notion of the revenue equivalence theorem is you can set s lower bar's value to whatever you want and then it comes out. And then you argue, well, in this case, we just keep lowering it.

GLENN ELLISON: Well, so I guess if you're thinking about optimal. If you're thinking about optimal, you could design mechanisms with the s lower bar type gets positive surplus. And then you wouldn't want to use-- those mechanisms would provide less revenue to the seller and more surplus to the buyer. So we tend to think of these as the mechanisms you would normally think someone would choose. And if you choose a mechanism to make sure the s lower bar type gets 0 surplus, then these things are going to hold.

Yeah, there are worse mechanisms that give positive surplus, the s lower bar types. There are better mechanisms that give negative surplus to the s lower bar types. But then how do you get the s lower types to participate in them? So in some sense, this theorem says that-- and sometimes I really don't like it, because it makes you think the field of auction design is completely irrelevant, because it says that all auctions work out the same.

And the reason why it's a-- let me see. So I have a proof here, again. Alex does a proof in [14.]124 if you want to see one. There's a proof. Actually, before I get to comments, let me just do one application of the revenue equivalence theorem.

Revenue equivalence theorem can also be a tool for solving auctions. So here's an example, an all pay auction. In the all pay auction, bidders submit bids b1 through bn. The high value bidder is awarded the object, but everyone pays what they bid, win or lose. So in this case, it's kind of like buying lottery tickets or raffle tickets. We just all submit our bids. Everyone actually has to pay their bid, win or lose. And then whoever bid the highest amount gets the object.

You can think of it as there are examples like lobbying. Imagine you're trying to win the rights to hold the next world cup in Qatar. Every country who wants a world cup submits bribes to FIFA. Whoever submits the highest bribe is going to win. If you submit the second highest bribe, you don't get your bribe back. You've given the millions of dollars to whatever person. We all pay our bribes and then the highest bribe wins.

This, again, would be an auction that the revenue equivalence theorem would apply to. And so because the revenue equivalence theorem applies, we know that a bidder of type si gets this much surplus. Integral of s lower bar to s f to the n minus 1 of x dx. But in this all pay auction, it's also an easy expression for your expected utility. Your expected utility in equilibrium is just your value times the probability that wealth has a lower value minus what you bid.

So if this minus this equals this, we get a markdown rule for your bid. Your bid is the expected surplus you're going to get from winning when you have the highest value minus your information rent. And so the revenue equivalence theorem, once you know it, can be a tool for actually computing equilibria of other auctions, because you know that the equilibrium must be such that people get exactly that surplus. And sometimes you can figure out. This case, you can figure out what the bids must be for everyone to get that surplus.

So anyway, it's a beautiful theorem. In practice, it's really not a good-- it has that basic thing to think about. We have to think about people bidding strategically and markdowns. There is an enormous field of auction design despite this. And that's because the revenue equivalence theorem really doesn't cover the interesting cases.

Many of the assumptions in revenue equivalence theorem are just not going to hold in practice. One is the bidders are assumed to be risk neutral. If you're talking about bidders bidding on paving contracts, they have businesses to run, they have employees who are going to be idle, it's not a long run thing. They can't be completely risk neutral.

Third one, values are independent. Even if I was having a hard time, I always have a hard time thinking of cases where the signals are actually going to be independent. So it could be that I actually know exactly how much I like this bottle of wine. I just have a valuation for it. I'm bidding for this value of wine against somebody else.

But in practice, if somebody else looked at that bottle of wine and they like it a lot, it's probably they know something about how good it is. If the wine actually tastes good, their signal is going to tell me something about how much I'm going to enjoy that bottle of wine.

If it's a painting I intend to hang on my wall, I know how much I like the painting. But the fact that someone else is willing to pay a lot for the painting, I always have the option of selling it to them later. So if there are a lot of other people who like this painting, then that painting is more valuable to me because of the option value of selling it.

Many, many cases, if I know that other-- if this is a good object, we're all going to it a lot. If this is a bad object, we're not going to it a lot. And so both the values are independent and the values only depend on my signal rather than on other people's signals as well. Those two things just don't seem realistic for many applications, because the fact that other people value something tells you something inherent about the object itself that then tells you something about my value. Or even just that I like it means that this is a bottle of wine that I unusually like. Someone else probably has the same tastes.

And then the second limitation is the set of mechanisms that it considers. It only considers mechanisms where you always allocate the good to the high value bidder. It doesn't let you not allocate the good. So just some examples of things people in IO have noted about how you would bid differently.

So suppose bidders are risk averse. They get utility u of s minus p if they win and u of minus p if they lose. So this is our small construction company that cares about its revenues in a nonlinear way. In this game, revenue equivalence theorem doesn't hold anymore and the first price auction raises more revenue than a second price auction.

Why would that be? Imagine you're in a first price auction and compare the equilibrium of the first price auction where you bid n minus 1 over n times v to the equivalent of the second price auction where you bid v. Those two auctions have exactly the same expected surplus in dollars to the bidder. That was the revenue equivalence theorem said.

But if in a first price auction, I raise my bid a little bit, I'm increasing the probability that I win, which is reducing the probability of the bad state where I get 0, and it's lowering my surplus in the good states where I do win. So what it's doing is it's to first order, it's not changing your expected payment by the revenue equivalence theorem, but it's reducing the variance in your payment by increasing the probability of winning a small amount instead of a lower probability of winning a bigger amount.

And so in the first price sealed bid auction, if everyone else was bidding b star of si, you would raise your bid in the first price sealed bid auction. Second price sealed bid auction is unchanged. You always still have the same dominant strategy. So with risk averse bidders, first price auctions raise more revenue than second price auctions.

Second example is if you have auctions with positively correlated values in auctions where the bidder's values are positively correlated, second price auctions raise more expected revenue than first price auctions. And there is a big literature on auction design and a big set of people who are professional auction designers, because to the extent that every auction differs from independent private values in some direction, you're going to think about all of these results that have been shown about when first and second or all pay auctions do better and pick out an auction that works well in that circumstance.

A second stark example of what this revenue equivalent theorem doesn't do is think about the case where there's n equals 1. There's one bidder. That bidder has a value that's uniformly distributed on 0, 1. This is the classic monopoly pricing problem that I started in lecture one with.

What does the revenue equivalence theorem consider? It considers mechanisms that always sell the good well. So if you have to always sell the good and the bidder has a value v that's drawn uniform 0, 1, the only mechanism that always sells the good is to sell it for a price of 0. And I started this class with a thing to do in this model is to set a price of 1/2 rather than a price of 0, because that maximizes probability of selling times value if you do sell. But if you have to sell with probability 1, all you can do is give the object away for free.

So revenue equivalence theorem mechanisms always give the object away for free to the s lower bar type. So you can charge s lower bar, but they always give away the good with probability 1, which means you're always giving away that surplus that we were in monopoly pricing where I was telling you, you don't want to give away. So revenue equivalence theorem is considering a limited class of mechanisms which don't include the mechanisms you would actually want to use.

So let me just think about reserve prices. What's your intuition for if we now go into an auction, do I want to set a reserve price in my auction or do I not want to set a reserve price in my auction? Or with one bidder, you know you need a reserve price. You want to set it equal to 1/2. What would I do with n bidders?

Two intuitions you might have is that the more bidders there are, the less need there is to set a reserve price, because the competition is going to set the reserve equivalent for me. But then on the alternate side, the more bidders there are, the less chance that the reserve price is binding. So the downside of the reserve, they don't sell the good also goes away.

So it turns out the theorem is in a second price sealed bid auction, you always set the monopoly price as the reserve. So what you want to do is, regardless of how many bidders there are, you just set a reserve price equal to the monopoly price. And I say that just before these bids are drawn, I just think ex-ante what would the monopoly price be to a single bidder? I announce in advance, this is the reserve. If you want the good, you must be at at least that reserve. You can't get the good unless you bid at least that reserve or the sale price will always be at least the reserve if you're the only one who bids above the reserve price.

And simple reasoning for that is think about a relaxed problem where the seller actually gets to choose a reserve price not in advance before conducting the auction, but after observing where the second highest bidder dropped out. So you had n bidders. n minus 1 of them have dropped out. And suppose you can cheat and after the n minus first bidder has dropped out. so someone dropped out at 2.3 million. I now know that this guy is willing to pay at least 2.3 million. And I think, what's the reserve price just for him? So it's a customized reserve price.

Well, so I have these n draws from this distribution. I know that n minus 1 of the draws are down here below this value, this drop out point, sn minus 1 colon n. What's my posterior on the value of the remaining bidder? My posterior on the value of the remaining bidder is just that it's a single draw from the upper tail of the distribution I started from.

So if I set a reserve price of r, which is bigger than this, the probability that I'm going to sell the good is just 1 minus f of r over 1 minus f of sn minus 1 colon n. It's just a single draw from the distribution conditional on the value being at least this. And so this is the inverse CDF. This is the upper, upper tail CDF of that.

So the expected profit that I'm going to get is just r times 1 minus f of r divided by this, which is just the constant. So if you think about how do I maximize this expression r times 1 minus f of r times this constant, I just maximize that by maximizing r times 1 minus f of r, which is the monopoly pricing problem for one customer.

So I have this monopoly profit when I use a price of r. Monopoly profits like this. If sn minus 1 colon n is here, if everyone else has dropped out low, I'm below the ex-ante monopoly price, then I would want to use the monopoly price as a reserve, because it's going to raise more revenue. Whereas if sn minus 1 colon n was up here, if everybody's bid really high, well above the monopoly price, then I don't want to use the reserve price anymore, because anything reserve price I do is just decreasing my revenue.

But in either case, just setting this in the low case, this is the binding reserve you want. In the high case, you don't want to reserve. And this is a non-binding reserve. So ex-post, it's always optimal to set pm as the reserve. And so ex-ante, it's always optimal to set pm as the reserve. So what you do in a second price sealed bid auction is you monopoly price as a reserve, and then that is the fully optimal mechanism.

So obviously if you're thinking about auction design, you need to think about what is that expected ex-ante monopoly price, because as the number of bidders goes up, the value of reserve prices do go down, because they're probably not going to be relevant anyway. But you still want the high reserve price. And competition doesn't mean don't use reserve prices.

Some comments here. Reserve prices are positive even if the seller doesn't have any value. Reserve prices do reduce social welfare. So this is an inefficiency of auction mechanisms is you use reserve prices, they create deadweight loss. Now, they only create deadweight loss if all the bidders have values below the reserve price. So it's less of an issue, but it's still a social loss.

The above analysis ignores entry. If you think about bidders like in the paving contract, paving Memorial Drive application, it's going to take bidders quite a long time to put together their bids and figure out exactly how much is it going to cost them to do the job. They're going to exert a lot of effort to submit a bid. If you put in a high reserve price, then they know that their surplus is going to be low. That's going to make you want to spend less time entering the auction.

Entry costs can make reserve prices not as good of an idea as you think, because if you put in a high entry, if there's a high reserve price and therefore people don't enter the auction, then attracting fewer bidders is going to be a cost. There's a theorem in a Bulow Klemperer paper showing that even in ex-ante sense, if you attract one fewer bidder in expectation, then you're actually worse off from having used the optimal reserve price.

And then another common reserve prices is there can be commitment problems in setting reserve price. We say we're only going to pay at most $2.1 million to repave Memorial Drive because it's got all these potholes in it. Everyone bids that it's going to take more than $2.1 million to do the work. What do you do then? Or if it's a bridge that's falling down, we're only going to pay $2.1 million to replace this bridge.

If everyone else bids more than that, what do you do? Do you really commit, OK, I'm never going to fix that bridge or repave that road? You're going to want to then put in a bid, say, OK, I guess $2.1 million was too low. Now I'll have a reserve of $2.3 million. Reserve prices only work if you can commit to them. They don't work if you undo them later on.

So then the final theory topic I'm going to do, I'll skip the next one after it, is common value auctions. More general model of auctions is that your value depends on your signal and everybody else's signal. So imagine you have these signals that have joint density f and your expected value depends on your signal and everybody else's signal.

A special case of this is the pure common values model. So suppose the object has value v to all bidders. And we get signals which are v plus epsilon. So I'm going to talk about an empirical paper about oil drilling. So imagine you're bidding for the right to drill for oil on some tract of land offshore. The value of winning that auction is the amount of oil that you're going to get to extract from that well.

And we all hire geologists and all of our geologists do sonic testing or whatever geologists do. And then they try to estimate what's the probability that there's oil. If there is oil, how much is there? And come up with an estimate of the value. And suppose they all get these values v plus epsilon i. The bidder's expected value given all signals is what's my expected value given my signal and what everybody else's geologists said. But I only observe my own signal.

So imagine we have this pure common values auction where there's this common value that everyone has for winning. We all submit our bids. And then I find out that I've won the bidding. When I find out that I've won the bidding, what I'm learning is that my geologist was more optimistic than everybody else's geologists.

And so particularly if there are n other geologists and n is a large number, I've learned that my guy was out in the tails in how he read the seismograph, which then tells me that this is what people call the winner's curse, that once I've learned that I win the auction, I now think the good is worth much less than I thought when I was submitting my bid, because I've learned that everybody else's estimate was lower than mine.

Now, in a rational model, it's not like bidders get negative value because of the winner's curse. In a rational model, what bidders do is they account for this in the winning. They're like, OK, I think there's $50 million worth of oil there. But what would I think if I learned that everybody else's geologists thought there was less oil than mine thought there was? And so what you bid is you bid based on what you think the value is conditional on learning that everybody else's geologist was less optimistic than yours.

So what you bid on is you're not bidding based on your-- so in a second price sealed bid auction, I'm not bidding what's my expected value conditional on my signal. I'm bidding what's my expected value conditional on my signal and on sj less than or equal to si for all j. And actually in terms of sj equal si for some j, but it's more complicated.

But anyway, what I'm bidding on, I'm bidding on what do I think it's worth knowing that everybody else is bidding has a worse signal than I do. And actually knowing that it actually matters what I bid. So there's somebody else who's very close to me. So what rational players will do is they avoid the winner's curse by just down weighting their bids based on the amount of winner's curse that they expect to occur once they learn that they won.

So what I want to do is finish today with one empirical paper. So this is about winner's curses in auction. So paper, Hendricks and Porter, empirical study of an auction with asymmetric information. This is about US auctions of leases to drill for oil in waters off the US Coast. Most of this oil drilling that they're talking about is if you go to the Gulf Coast of the United States, off the Coast of Texas, Louisiana, there are many, many, many offshore oil wells and in relatively shallow water where the government has auctioned off the right to drill for oil.

The US government auctions off two different types of offshore oil leases. One are what are called wildcat tracts. Wildcat tracts are roughly 5,000 acre plots that don't abut any other drilling. And about 35% of those wildcat tracts that are bought end up producing any oil. The others, either the company doesn't drill them because doing a more detailed geological survey after they've won the auction, they decide it's not even worth drilling. Or they drill a hole and they spend millions of dollars drilling underground and then they don't find oil when they drill.

The second type of things that are auctioned is that I'll just go out here in the Gulf of Mexico and they'll just auction off the right to drill here. But then sometimes what happens is you auction off the right to drill on this tract and there is oil there. And if there is oil there, then immediately people tend to want to drill here or here or here or here or here. Because if you know there's oil here, there's probably oil nearby also.

And so they also auction off what they call drainage tracts. They're smaller plots that are adjacent to a previously drilled wildcat tract. And so when people buy drainage tracts, it's more likely that there's going to be oil there. It's much less risky, much higher expected amount of oil, because you know there's oil right next door.

Here's just a map showing-- here's random wildcat tract that someone is planning to drill on. This would be an area where, say, this tract was drilled. They had oil. And then they started auctioning off all these little drainage tracts around it where people can think if there's oil here, there's probably oil here and all these other places nearby.

And asymmetric information that the winner's curse, they note, is that this auction is very asymmetric auction in that there's one firm that has a well in the ground at this location that knows something about the oil that they're pumping out of it. And then there are all these people around here who don't have the same information that this person with the wildcat tract is here.

And particularly, the way my understanding of oil drilling is, oil drilling is a fairly sophisticated thing. So you've got this thing that's underwater. I've got my well here that I'm drilling down. I'm going underwater, and then I'm drilling into the ground on my wildcat. So I own this wildcat tract. You can actually send a drill bit diagonally underground. So you don't have to drill a straight hole. You can drill like this. You can drill like this.

And so imagine that this firm is here. They're producing a ton of oil from a well that you know on the surface is up there. And then they propose that there's an auction going to be conducted not only for this wildcat tract, but for this. So this is the wildcat tract. There's a drainage tract here. There's a drainage tract here. And you're going to bid against ExxonMobil, who's got their pipe in the ground, pumping a ton of oil out of this wildcat tract.

And let's suppose that you're bidding against them and you're like, OK, I'm going to bid for this drainage tract. And what you learn is that they didn't bid for that drainage tract. They bid for that drainage tract. As soon as you learn that ExxonMobil bid for this drainage tract, you think it must be that what they did is they drilled a thing down here and they went diagonally that way. So they think there's oil over here. I've just won this drainage tract. Do I even want to drill it once I've learned that ExxonMobil didn't want this one?

Anyway, so that creates the potential for a severe winner's curse and that the person who has the pipe in the ground knows much more about the oil pressure and about where the oil is coming from than all the firms who are just able to observe that they've got a rig set up here in the middle of the ocean somewhere.

One thing that's interesting about this is so these drainage tracts, if you look at the fraction that are productive, and I guess it's even really hard to see this, so there were 1,000 in this data set. There were 1,056 wildcat tracks. Only 385 had oil. There were 144 drainage tracks. 86 had oil. So it's much higher fraction have oil. The average value of the oil taken out of them, these tracts have an average of 5.27-- I guess it must be millions of something or other dollars in oil. The drainage tracts have $13.5 million of value on average.

Yet somehow the wildcat tracts attract 3 and 1/2 bidders per auction. The drainage tracts attract only 2.7 bidders per auction. So you have these much more valuable tracts that are much lower risk and yet attract fewer bidders than the wildcat tracts. And as a result, unsurprisingly, the average profits earned by the winning bidder are much higher in bidding for drainage tracts than bidding for wildcat tracts.

And Hendricks and Porter started to note, isn't this initially seem like a puzzle? Why would the drainage tracts that are so much more reliable and so much more profitable have fewer bidders and much higher profits than the wildcat ones where it's highly risky? And they suggest that this is a natural consequence of the winner's curse. That the because the winner's curse is so severe, everybody is afraid to bid against the incumbent who's already extracting the oil from the wildcat tract. And therefore, the incumbents earn high profits on these tracts and other people just tend to stay away because they're afraid of the winner's curse.

So it's a nice idea. But in some ways, I think I wanted to cover Hendricks and Porter is it's a really beautiful example of how you can write a reduced form empirical IO paper in that they start with of simple idea of people are afraid of the winner's curse, therefore they bid less. And what they try to do is build that out as-- write down a model and then think of seven different implications of the model that are not obvious from when you just state the model, and then think about are those implications all true.

And use that to decide, is the asymmetric information really the explanation here, or could there be some different explanation? They don't work out what the alternate explanations are, but the thought is if I have my one explanation, my one explanation makes seven different predictions. I can then test my seven different predictions, if they seem to hold, then I think this is a good explanation for this.

So here's a very simple model. So they have a wildcat tract that has one informed person who's drilling on the wildcat tract, and then there's a single drainage tract nearby. They're going to assume that the drainage tract has value v to the person who is already drilling the wildcat tract and v minus c to the others. So the option is not only does ExxonMobil have an informational advantage.

ExxonMobil would get v if they win the drainage tract, whereas a rival would only get v minus c. And the thought is that the rivals are getting less value because when you already have a platform nearby, there's economies of scale and building a second platform here. And so there's a cost advantage as well that might explain why ExxonMobil will bid more.

We have n uninformed bidders. The uninformed bidders do get some signal. They get some signal z, which can be the signal from their geologists. It could be the signal from ExxonMobil does have to report how much oil they're pulling out of that well because they pay royalties. And they're going to assume that this satisfies the expected value is bigger than the government's reserve price. The expected value is increasing in the quality of their signal. But there's always some probability that the value will be less than the reserve price, no matter how positive the signal is.

And then you have the one informed firm that both sees the true value, actually fully knows how much oil is under the drainage tract, and sees the signal that the uninformed bidders are relying on. And then they think about a first price sealed bid auction with a reservation price of r.

Observation like we've done a few times in this class. This is not a game that can have a pure strategy equilibrium. Because in a pure strategy equilibrium in some sense, the firm playing the uninformed guys are just too vulnerable when they play pure strategy equilibrium. So you can't have an equilibrium where the uninformed bidders are not bidding, because the uninformed guys were not bidding, the informed guy would just bid r.

And then bidding r plus epsilon would be a profitable deviation for the uninformed bidders, because on average, the tracts are worth a lot. So you can't have the uninformed bidders not bidding, because then the informed bidder would bid r, and then the uninformed guys would wish they were bidding r plus epsilon. And so that can't be an equilibrium.

You also can't have an equilibrium where the uninformed guys are following some pure strategy that's some function b uninformed of z. Because if they were playing this pure strategy, then what would the incumbent do? If the value was bigger than what they're bidding, the incumbent would take it. And if the value was less than what they're bidding, the incumbent will leave it to them. So basically, if they're playing a pure strategy, they always lose. Because either the incumbent, it's a good tract and the incumbent takes it from them, or it's a bad tract and the incumbent leaves it with them.

So you can't have any of these equilibria. What do you get in these models of search? You get a mixed strategy equilibria. So the uninformed bidders are going to mix. Sometimes they're going to bid exactly 0 and sometimes they're going to mix over the interval from r to z. And their mixing is going to be such that it forces the incumbent to play. It's optimal for the incumbent to play a strategy, which is if there's no oil bid 0, if v is a little bit bigger than the reserve price and z is not too big, then just bid exactly the reserve price. And then as either the value gets larger or z gets bigger, start to increase your bid and bid above the reserve price.

So the incumbent is, if I think about what the incumbent's bid distribution looks like, sometimes I see the incumbent bidding exactly 0. Sorry, bidding exactly 0. Just not even bidding. Sometimes I see the incumbent bidding the reserve price. And then there's going to be some density of bids where the incumbent bids more if it knows the tract is highly valuable and looks attractive. It's going to bid a higher amount.

So this is what the incumbent does. What do the uninformed guys do? The uninformed guys are playing a corresponding mixture. They don't know anything, but sometimes they bid 0 and then they also bid over some distribution of higher values. Let me be more realistic. Maybe they bid 0 a lot because they're scared.

So anyway, Hendricks and Porter solve for that equilibrium. And a very nice feature is that they have many different predictions that they can derive in this model. So the first one is that in this equilibrium, the incumbents bid 0 with a smaller probability than the probability that the uninformed guys bid 0. So 0 bids by the incumbent are less common than 0 bids by the disadvantaged entrants.

The informed bidder wins the auction at least half the time. The uninformed bidders are earning 0 profits because they're mixing and they're mixing with 0 in the equilibrium support. So sometimes they do bid 0. And so profits are 0 in expectation. In the bad news case where they bid something and the informed bidder bids 0, then their profits are negative. And in the good news case, where the informed bidder bids at least r, then their profits are positive in expectation, because they have 0 expected profits. In the bad news case, they get negative profits. They must get positive profits in the good news case.

The incumbent earns positive profits. A much less obvious statement is that when the cost c is close to 0, so the two firms don't have big cost differences, the ex-ante bid distributions of the incumbents and the entrants and the uninformed bidders converge to each other. So as c goes to 0, this is a bid distribution that comes from bidding based on what you know about the value. This is one that comes purely from randomization. But these two distributions approach each other as c goes to 0. And the c equals 0 limit, these two distributions would be identical.

The incumbent's bid is independent of the number of rival bids. That's another property of the mixing. And the incumbent's bid is an increasing function of the true value of the tract and of the signal that's available when people are bidding on the tract. And it's an increasing function of the true value, because the more it's worth winning, the more aggressive you want to be. It's an increasing function of z, because if you know the opponents are getting better signals, they're more likely to enter, you bid more aggressively.

And then what Hendricks and Porter do is they get this data set. It's a somewhat historical paper. They got information on 114 drainage tracts auctioned between 1959 and 1969. All of them are adjacent to a wildcat tract on which they have production data. And so one of the nice things about going back historically is they're doing this far enough after the time all these auctions occurred that they can observe all the production that occurred over the lifetime of the wells.

So the wells are being drilled over the next five years after their auction, and then they're producing oil after some number of years, but they're able to run out when all the oil got produced and how valuable things were ex post. Anyway, the data include the bids by all the participants, the number of neighbors bidding, the number of non-neighbors bidding, the ex post profitability of the tract, the gross profits of the adjacent tract, and things like that.

One of the interesting design features the US government chose is imagine sometimes you may have this situation. So imagine you have ExxonMobil here and you have BP here. And then there's a drainage tract right here. The one case you would think, OK, then you're going to get a lot of competition. If ExxonMobil's got this, BP's got this, they're both producing oil and there's this little area in between.

Rules the US government set up for this is that ExxonMobil and BP in that situation can form a joint venture and bid together jointly on the tract rather than competing with each other. So in most cases, in the data set where this occurs, lo and behold, the firms do form the joint venture and bid together. So mostly they have one incumbent tract.

Anyway, so then testing predictions of the model. Prediction one that the neighbors are supposed to bid much higher fraction time than the non-neighbors. That's borne out in the data. The neighbors are bidding 83%, the non-neighbors, 68%. The incumbents are supposed to win more than half the time. They're winning 52% of the time. I don't know if that's such great confirmation.

The uninformed bidders are supposed to get negative profits if the incumbent enters, negative profits otherwise. That's one of the clearer things in the data that when auctions where the non neighbors win, you compare the situations where the incumbent bid and the incumbent didn't bid. When the incumbent didn't bid, you get very negative profits. When the incumbent did bid, you get positive profits.

The incumbents are earning a lot of money. That was also very clear in this table. The incumbent win tracts are very high positive profits, whereas these are essentially 0 on average.

The bottom part of the paper is where this starts to seem less high tech econometrically than you would expect today. We're trying to see whether these two distributions are ex-ante identical rather than doing some kind of non-parametric estimation of the densities of those distributions. They're just running some regressions and asking whether the epsilons seem to have similar variances and whether the bids seem to be similarly related to the number once you control for the acres, the number of neighbors, the values, things like that. So I don't really know whether that's really great verification.

Then my prediction six was that the incumbent's bid doesn't depend on the number of non-incumbents bidding. That's also held out in a regression. And then to check that the incumbent bid goes up in the v and the z, they show the incumbent bid does increase in the true value of oil that's eventually taken out of the tract 20 years later and also in the amount of money that comes out of the wildcat tract, which is the observable thing that everyone else can see.

So is that my last slide or do I have more? Nope, that's it. So anyway, I think it's a very, very nice example of how you think about this idea of writing reduced form IO where you just, here's my model, here's a bunch of predictions that are not completely obvious, and then I'm going to look at as many more nuanced predictions as I can, find that they're true. And it seems like a pretty convincing case that we do badly in these auctions because of the asymmetric information.

Can I say that-- and sometimes people have criticized the US government and saying, well, why on Earth would you give BP and ExxonMobil the ability to jointly bid here? One way to think about this is, though, maybe that's actually not so stupid. Maybe what we want to do on the wildcat tracts is in some sense it's fine to have adverse selection on the wildcat tracts, give all the money away to the person who has the initial-- sorry, give all the drainage tract profits away to the person who drilled the initial wildcat tracts, because the wildcat tracts, again, were out here in the middle of nowhere.

There it's a symmetric auction model. So in some sense, what you're doing is just increasing the prize you get from the symmetric auction model to be you get the prize, which is all the oil here, plus the advantage on the surrounding drainage tracts. And so in some sense, if we just always give you this big secondary prize, that just gets competed away in the bidding for the initial wildcat tracts. And we can use the competition for the wildcat tracts as the sensible way to try to raise revenue, not worrying so much about trying to get the hard to get drainage tract revenue from the firms and just bundle it in with the initial wildcat revenue and then let them bid symmetrically in that case.