[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: Today we're going to talk about something called revenue equivalence.

So so far, we've studied a few different auction formats, and we've solved for the equilibria of these auction formats. And we found this surprising or peculiar finding that the amount of revenue that these auctions generate doesn't seem to depend on which auction format we use. So I want to first review a few of those examples. And then, today, we're going to try to understand, in general, why this happens and what are the limitations and scope of this result.

So remember, we first looked at the really simple case where we had only two bidders, and each bidder's valuation was uniformly and independently drawn from the 0, 1 interval. So in the simple case, we had v1 and v2. This is bidder 1's valuation. This is bidder 2's valuation. They were in 0, 1, and they were uniform and independent.

And in this really simple case, we looked at the first-price auction, and we said that, in the first-price auction, there was an equilibrium. Maybe I'll write b first-price auction where bidder i's bid, if their valuation was vi, was half of their true valuation. We know, in the first-price auction, people bid in equilibrium less than their true valuation.

This is sometimes called bid shading. You shade your bid below your valuation. And we saw that this was exactly vi over 2.

And then we also found an equilibrium of the Second-Price Auction, so I'll write SPA. And in the second-price auction, we observed that the amount you pay for the item, if you win, doesn't depend on your bid. It only depends on other people's bids. And we gave this argument as to why you should always bid truthfully in a second-price auction. So here, we had bi of vi equals vi.

But what we noticed is, well, let's try to compute in each case how much you are expected to pay in each of these auctions. Well, in a first-price auction, let's say your true valuation is vi and this is how much you bid, well, your expected payment is going to be your bid times the probability that you actually win the good because you only pay your bid when you win. So your expected payment is vi over 2. That's your bid.

And in a first-price auction, you pay your bid if you win. So it's the amount you pay times the probability that you win. What is the probability that you win in this auction if your valuation is vi?

So maybe I'll say this is going to be the probability of winning. Well, this was a symmetric equilibrium. So the bidder who wins is the bidder who has the highest valuation. So I'm going to win-- yeah.

AUDIENCE: v1, right?

IAN BALL: Exactly, and I guess here we have vi, so it's going to be vi, but exactly. Yeah, it could be 1. Exactly right.

Why is it vi? Well, I'm going to win if my value is higher than everyone else's because we're all using the same increasing bidding function. What's the probability that my value is higher than everyone else's? Well, there's only two of us, so it's the probability that the other bidder has a value below mine.

But for a uniform distribution, if my value is vi, the probability that the other bidder's value is below vi is just vi. This is just, for uniform, 0, 1. Here's my bid. Here's my value vi.

I'm going to win if the other bidders value lies in this interval, and the length of this interval is exactly vi. So we multiply it by the probability of winning, and we get vi squared over 2. This is the amount that I expect to pay if my valuation is vi.

And then we can do a similar calculation here. What is my expected payment in this case? Well, we're still going to multiply by vi, the probability of winning.

But now we can't just multiply my probability of winning times the amount I pay if I win because the amount that I pay if I win depends on what other people bid. So what we have to multiply it by, is not what I pay if I win, but my expected payment conditional on winning.

And this is a general way of breaking down how much I expect to pay. I only pay if I win, so the expected amount that I pay is my probability of winning times my expected payment if I win. My expected payment conditional on winning.

Now, if I win, how much do I pay? Will I pay the second highest bid? If I win, the second highest bid is the bit of the other person, and I know the other person is bidding truthfully.

So the amount that I'm going to pay is my opponent's valuation. And it turns out what this is going to be is, well, let's see. It's the expectation of maybe vj conditional on vj being less than or equal to vi.

So I'm bidder i. I know that my valuation is vi. When do I win? I win exactly if my opponent's value vj is less than mine. And if I do win, I pay exactly their valuation because I pay their bid and their bid equals their valuation.

So we're going to get this formula here. Maybe I'll erase this arrow just to give myself some more space. So this is my probability of winning.

This is my expected payment conditional on winning. And can we compute this conditional expectation here? Yes.

AUDIENCE: Yes, we can compute it. And actually, we can also think about how vj is uniformly distributed from 0 to vi so that the expectation would just be half of vi.

IAN BALL: Exactly, let's be careful. So vj is uniformly distributed from 0 to 1. I mean, you're right, but you're saying conditional on vj being less than or equal to vi, it's uniformly distributed on 0 to vi.

And therefore, here's the example, if I know that the other player's valuation is below mine, well, I know it's somewhere in here. And if it's uniform, Well, it turns out, conditionally, it's still uniform over here. What's its mean?

Well, it's just going to be the middle of this interval, which is exactly vi over 2. So exactly here, vi over 2. So if I compute this out, what do I get? I also get vi over 2 times vi, and what do I get? vi squared over 2.

So here, this is just a recap of this finding we had before that, in both the first-price and the second-price auction, at least in the equilibrium that we looked at, the expected payment made by bidder i, when her valuation is vi, is exactly the same in the two auctions. And that seemed kind of puzzling. And you wonder if that's a coincidence.

Well, more generally, let's look at the general case. So in the general case, the general case, we had n bidders instead of just 2, and we didn't assume that they were uniformly distributed. We allowed more general distribution.

So here, we had the valuations were v1 up to vn. And I think what we said is they lived in some set v lower bar to v upper bar where everything is nonnegative. So everyone's value lies somewhere between this lower bound v lower bar and this upper bound v upper bar. Up here, those bounds were just 0 and 1, but here we're going to be more general.

And we're going to assume that these valuations are independent. That is, the distribution of v2 or v3 is statistically independent of the distribution of v1. But now they're drawn from a different distribution. Instead of the uniform, we said they have density f over this interval, the lower bar to the upper bar.

And then we introduced one more piece of notation. It was nice to have notation for the distribution of the maximum of everyone else's valuation. So we looked at the max over, let's say, j not equal to i of vj.

So I'm holding out one bidder. I'm taking the perspective of bidder i, and I'm thinking of all the bidders other than bidder i. I'm looking at all of their valuations, and I'm taking the maximum.

This is another random variable. As bidder i, I don't know what this is. I can form beliefs about this. And what we said is we just had notation.

We said this has cumulative distribution function big G and density little g. And we had formulas for what these are, but the exact formulas don't matter. It's just that we can compute the distribution of this. And we gave a name to it. The cumulative distribution is big G, and the density is little g.

Notice here that everything here is symmetric. So the value of i here doesn't matter. If I'm bidder 1 and I'm thinking about, what's the highest value of everyone other than me? Or if I'm bidder 2 and I'm saying, what's the highest value of everyone other than me? That distribution is going to be the same.

Maybe just make sure we're on the same page. If n equals 2, what is the density g going to be? If n equals 2, so there's two bidders. Yeah.

AUDIENCE: f.

IAN BALL: Just f, because if there's two bidders, well, what is the maximum of every other bidder except me? Well, it's just the other person. So this is just vj.

And we already assumed that the distribution of each bidder's valuation follows density f. So in this case, g is just going to reduce to f because there's only one other bidder. But when n is bigger than 2, then this is going to be a more complicated form. But the simple thing is the highest bidder among my opponents, if I only have one opponent, is just the value of my opponent. OK, great.

So here was the general case, and we also looked at the first-price auction and the second-price auction in this case. And we computed equilibria. And maybe I'll start here with a second-price auction because that was-- maybe that's a bit easier to write down. So bi vi, and the second price-auction-- people remember what our equilibrium was here?

I think people know. They just don't want to say. If we have a second-price auction, how much should you bid if your valuation is vi? Yeah.

AUDIENCE: Vi

IAN BALL: Just Vi, you should just bid your valuation. So it turned out, we gave this argument that, in a second-price auction, you always want to bid truthfully. You always want to bid your valuation, exactly. And that didn't depend on the details of whether there's two bidders or n bidders or what the distributions are. That's a very robust result.

So we've got vi. And we can, again, try to compute, what is your expected payment? OK, this is a little trickier here, but let's see if we can do it.

Well, first, let's just write it abstractly. What is my expected payment? Well, we first have to say, well, there's some chance I don't pay anything at all.

If I lose the auction, I pay nothing. So my expected payment, it's always helpful to break it down into the probability that I win the auction times the expected amount that I pay, given that I win. So we'll break it down just like we did before.

So let me first write down the probability that I win. So this is, the probability that I win is the probability-- well, in this equilibrium, everyone's bidding their valuation. So I'm going to win if my bid is highest, which is exactly the case if my value is higher than everyone else's.

So I'm going to win if the probability of the max of j not equal to i of vj is less than or equal to vi. And as usual, we don't have to be too careful. Is it less than or equal?

Is it strictly less than. The probability is going to be the same, so we don't really have to-- it doesn't really matter what we write here. So this is the probability that I win.

Now what is the expected payment conditional on winning? So it's going to be an analog of this formula here. The expected payment conditional on winning, well it's going to be an expectation-- a little more space here. Let's see how small I can write.

Well, it's expectation conditional on me winning. So I'm exactly going to win if the max over j not equal to i vj is less than or equal to vi. This is exactly the event that I win. I'm going to win if everyone else's valuation is below mine. Another way of saying that is the maximum over everyone else's valuation is weakly smaller than mine.

And when that happens, how much am I going to pay in the second-price auction? Yeah.

AUDIENCE: vj. Like, the second highest.

IAN BALL: Right, and we have some notation for that right here. It's exactly-- so you're exactly right. I'm going to pay the second highest valuation.

But what that is is it's precisely the highest among everyone other than me because I'm the highest. So you're exactly right. It is-- let's write it as this, max vj not equal to i of vj.

So it's crucial here, yes, that this is the second highest value. But in the event that I have the highest value, well, what's the second highest value? It's just the highest among everyone else. So we exactly get this formula here.

And this gets maybe a little messy, but we're actually able to compute it. Let's see. So it's going to be an integral from v lower bar to vi.

So technically, we get into some conditional densities. It gets a little messy, but let me just write out what we get. We're going to get t-- maybe I'll write x, g of x for g of vi dx.

So I just wrote this probability as g of vi because big G is the cumulative distribution function associated with this random variable. So the probability that this random variable is below vi is, by definition, g of vi. So here, I'm just using a definition.

For this part, it's a little trickier. I'm using what's called a conditional density. I'm saying, what is the conditional density of the highest value of my opponents, given that their valuation is below vi?

And this here is my conditional density. I'm integrating this between v lower bar and vi because I'm focusing on the event that the highest value of my opponents is lower than mine. And maybe we don't need to get too caught up in the algebra here, but we can just see that this G of vi and this G of vi cancel, and we're going to get a really simple formula.

So this is my general formula for how much I expect to pay in the second price auction if my valuation is vi. Let's do one more. Let's finish things off here.

And as I said, if this conditional density is new to you, it won't come up again. So the final answer is maybe what's important here.

Let's just do one more, the bi first-price auction of vi. We actually computed this. Anyone remember what this was? We computed this in class before. Maybe I'll just write it down.

It turns out we had a formula that actually looked exactly like this. This was actually our formula. So it turns out we derived this in class, I think, last Thursday. We used a differential equation, and we found that it's exactly this.

So this is just a formula that we had. But now if this is how much I'm bidding in a first-price auction, what is the expected payment that I'm going to make? Well, in first-price auction, what I pay is my bid?

So my expected payment is just my bid, how much I pay, times the probability that I actually have to pay my bid. But when do I pay my bid? I pay my bid when I win.

So my expected payment, well, it's just bi FPA of vi, my bid times the probability that I win. What is the probability that I win? Well, it's the probability that my value is higher than everyone else. It's exactly this thing we already computed down here, and we have a simple formula for it. It's just g of vi.

And this is probability of winning. OK, well, we already have a formula for bi FPA of vi, so let's plug in this here. And we see that the G of vi cancels, and we get the same integral.

And we get this. So this seems like it probably is not a coincidence, that we've gotten exactly the same answer in both cases. Yes.

AUDIENCE: So why is bi of vi this integral when we had already said it was just vi?

IAN BALL: The key is the superscript here, SPA. So in a second-price auction, the equilibrium bidding strategy is for me to bid exactly my valuation. So everything we did here was for the second-price auction.

Now we're changing cases here-- and I should've been more clear about this. Now we're moving to the first-price auction. In the first-price auction, the equilibrium is different. In the first price auction, I don't want to bid my true valuation.

If I bid my true valuation in a first-price auction, then I can never come out ahead because either I don't win, or I win, but then I have to pay my bid. So deriving this was a lot harder. We spent basically all of class deriving a differential equation. And in the end, we got this formula.

So if you don't-- this is hard to see. This took us an hour to get to, but it should be in your notes from last class. So the answer to, why are these different? These are different equilibria of different auctions.

But the surprising thing is we get exactly the same formula for the expected payment in both cases. And let's just see if that's consistent. Up here, I got vi squared over 2. Is this formula going to give me vi squared over 2 in the really simple example, in the simple uniform example? Can we see this little calculus exercise?

Well, in the uniform example, what is g? In our example, where we had two bidders, and the values were uniformly distributed over 0 1? It's just 1 because the first observation that you made was that d equals f when we have two bidders.

But if our valuations are uniformly distributed, what is f? It's just 1. So if we just plug in 1 here, we get the integral from v lower bar to vi.

Well, so let me go through our example. So in our simple example, everything works out because g of x equals 1. And what is v lower bar?

AUDIENCE: 0.

IAN BALL: 0, right? And v upper bar is 1, but that doesn't show up. So now we just get the integral from 0 to vi of xdx.

And this is something people should know. This is just the antiderivative. This is just x squared over 2, so we just get vi squared over 2.

And this just confirms the two formulas that we got up here. So everything's consistent. That's good.

And we have this kind of surprising finding. So note the equilibrium expected payment is the same in the first price auction and the second price auction.

To be a little more precise, I'm talking about a particular equilibrium that we computed. But I'll be a little loose here. The equilibrium expected payment is the same in the first-price auction and the second-price auction.

And notice how this happened came about, getting back to the question, the rules of these auctions are different. So if the players bid the same way in these two auctions, their payments would definitely be different. But the equilibrium bidding behavior is also different. So the rules of the auction are different, the bidding behavior is different, and somehow, these things offset in exactly the right way that we get exactly the same formula in both cases.

And I think it feels, to me at least, that this can't be a coincidence. It seems like there's something else going on here. And what's going on is exactly what's called revenue equivalence.

And today, we're going to try to understand more generally why and under what conditions this is true. So I think this should be kind of a puzzle. It's not at all clear why we get the same answer here. And today, we want to get to the bottom and understand this puzzle.

But this was to be motivation for what we do. So if this is unclear, what we do won't be motivated. So let me pause for some questions here. Everything clear so far? OK.

Well, so far, we said that the equilibrium expected payment was the same in these two auctions. What if we used a different auction format? Maybe we could use a third-price auction or a fourth-price auction, all-pay auction. There's many different auction formats we could use.

And we'd like our theorem to speak to any auction format you could come up with. So to do that, we have to think, well, what is an auction? How can I define a general abstract auction?

So let's be really general. What is-- maybe I'll call this a generalized auction. And I'm still going to be in this setting over here.

So I'm still assuming that these assumptions about the environment, these assumptions about how much each person values the item, but what's going to be general is the auction format that we use. Well, at bottom, what is an auction? Well, first, I have to specify what bids are allowed.

So far, we've allowed bids to be real numbers. But you can have really weird auctions where bids can only be multiples of 5. Or maybe you can only say yes or no. Or maybe you say I really like it, or I don't really like it. There's a lot of different ways we could do it.

So a generalized auction format, first, we have to specify the bid sets. B1, B2, all the way up to Bn. So this is specifying what each bidder is allowed to bid.

You could think of these as messages. Or there's a lot of different ways we can think about it, but this is just the set of choices that each bidder has in the auction. And then we need to specify an allocation rule.

An allocation rule says, well, if we collect some bid by bidder 1 and some bid by bidder 2 and all the way up to a bid by bidder n, what do we do? Well, the auction is to specify two things. The allocation rule is, who gets the good? And then the payment rule is, how much does everyone pay? So we'll call it the allocation rule and maybe the payment or transfer rule.

Now, for instance, so far, we've only looked at auctions where, all the auctions we've looked at, you only pay if you win. But that's actually not a requirement of an auction. In fact, there's something called an all-pay auction where everyone submits their bid, and you have to pay your bid whether or not you win the item or not.

That may sound like an unappealing auction, but of course, you're going to bid less in response to that. So we're actually going to allow pretty general payment rules and transfer rules. We've generally looked at allocation rules where the highest bidder gets the item.

But we could be even more general. Maybe we like one of the bidders more, and we want to always give them the item. So in general, what is an allocation rule?

It's going to be a function-- maybe I'll call it q-- from B1 plus Bn to, well, it's going to take in a profile of bids by all the players, and what it's going to spit out is it should spit out something like, I give the good to player 1, or I give the good to player 2, or I give the good to player n. The problem is sometimes you have to randomize because sometimes two bidders will bid the same amount, and then we might have to flip a coin to see who gets the item. So we have to allow these kinds of lotteries.

So I'm going to write-- it's going to map to what I'll call delta, where delta is just going to be the set of lotteries. It's going to be the set of all vectors q1 up to qn, where the interpretation is, if these are the submitted bids, then q1 is going to tell us the probability that player 1 gets the item. So these are going to be-- so q1 to qn is greater than or equal to 0.

Now, I want to emphasize this here. I don't require equality. So what's the interpretation here? What if I were to say, OK, if these are the bids that are submitted, I'm going to allocate the good with these probabilities that sum to strictly less than 1.

Could this ever happen? What's the interpretation? Yeah.

AUDIENCE: Well, there's also the situation in which no one gets it.

IAN BALL: Right, exactly, so why do I have to allocate the good? I might say-- and we see this in auctions-- no one's bid was high enough. I'm not happy with the bids people submitted.

I'm going to keep the good myself. And that would correspond to all of these being 0, actually. So we allow the auctioneer to keep the good.

So more formally, 1 minus the sum of the qi's is exactly the probability that the auctioneer keeps the good. So one way of thinking about this is it's a probability distribution, not over bidders 1 through n, but over the auctioneer and bidders 1 through n. And it doesn't have to sum up to 1 because the auctioneer could get the good with some probability.

So this is just a definition of this set delta that I'm going to use. So delta is the set of, sometimes it's called, subprobabilities. It's like distributions over who gets the good. But they might sum to less than 1.

Now what about a payment or transfer rule? That's going to be a function t to Rn where, as a function of the bids, I'm going to specify t1 how much player 1 has to pay me; T2, how much player 2 has to pay me; and so on.

So just to make sure we understand what's going on, let's look at a really simple example. Let's look at the first-price auction. So in the standard first-price auction, what were these bid sets, B1 through Bn? What did we allow people to bid? Yeah.

AUDIENCE: Any real number that's positive.

IAN BALL: Any real number that's positive or maybe nonnegative. Yeah, so I think it doesn't really-- the details don't matter too much if we allow negative bids. But you're right, so far, we haven't allowed negative bids.

So in that case, we had B1 is equal to Bn is equal to, maybe I'll write, R plus, or maybe I'll be a bit more explicit. I'll write 0 to infinity. So what we said is, what is the bid set for each player? The bid set is just the set of all nonnegative real numbers. That's what we allowed you to bid.

Maybe, in reality, maybe you can only bid multiples of a set or something, and we could add more complexity, but we'll just allow this for now. Now, specifying the allocation rule and transfer rule, it's just a bit messy to write it all out because their whole function. So let's just go through an example, one example.

What if, say, n equals 2? So we have two bidders. Let's see if we can compute q of 3, 4 and t of 3, 4, so really just to make sure we understand the notation that's going on.

So if n equals 2, the bid set for player 1 is just a nonnegative real number. The bid set for player 2 is just the set of nonnegative real numbers. So the domain of the allocation rule and the transfer rule, they have to say if bidder 1 bids this number and bidder 2 bids this number, who do we allocate the good to? And how much do we make them pay?

So when I write q of 3, 4, the interpretation of this is 3 is the bid by bidder 1. 4 is the bid by bidder 2. And we need to say what happens if bidder 1 bids 3 and bidder 2 bids 4?

And this is really just to check our notation here. So what is q of 3, 4? Yeah.

AUDIENCE: Would it be a 0 comma 1.

IAN BALL: 0 comma 1, right? Because here, because bidder 2's bid is higher than bidder 1's bid, we need to allocate the good to bidder 2. And we have notation for that. Our notation just tells us the probability that each bidder gets the good.

Now probabilities seem a bit weird here. They're just 0, 1. We always give the good to bidder 2. If bidder 2's bid is higher. So this is just going to be the vector 0 comma 1 that says we give the good to bidder 1 with probability 0, and we give the good to bidder 2 with probability 1.

What about the vector t, t of 3, 4. What is that going to be? Yeah.

AUDIENCE: 0, 4.

IAN BALL: 0, 4. This says, well, this first component we might call t1 of 3, 4. This is how much bidder 1 has to pay. Well, bidder 1 got outbid, they're not getting the item, and therefore they don't pay. So they pay 0.

Bidder 2 was the highest bidder. They're winning the item, and in the rules of the first-price auction, they pay their bid. So we get 0, 4.

Now, what if I were to change this to a second-price auction? If I change this to a second-price auction, what would change here? Would q change?

Would t change? What would be different? Yeah.

AUDIENCE: q would not change because, still, the person bidding 4-- well, yeah, they'll bid 4. In this case, would win. T would change, however, because they would only have to pay 3.

IAN BALL: Exactly, so notice, when we compare the first- and second-price auction, they actually use the same allocation rule. And a lot of auctions use this allocation rule. The allocation rule is pretty simple-- we allocate it to the highest bidder. They both use that allocation rule.

The difference between the first-price auction and the second-price auction is the payment rule, how much people have to pay. And indeed, down here, what we would get is 0, 3. If it were a second-price auction, the first bidder got outbid. They still lose the item and pay nothing.

The second bidder, in a second-price auction, though, pays the-- well, bidder 2 is the highest bidder. They pay the second highest bid. Well, what's the second highest bid between 3 and 4? It would just be 3. So we'd have to change this 4 to a 3 if we wanted to do the first-price auction.

What about an all-pay auction? What would this be if it were an all-pay auction? Well, then it would just be 3, 4 because bidder 1 would have to pay 3 even though they didn't when the good and bidder 2 would pay 4. They'd pay their own bid. OK, so I think we understand the notation here.

And now that we have this generalized auction format, we want to try to understand these expected payments. Now, to talk about how much people expect to pay, we have to first have an equilibrium in mind.

This is a key conceptual point that, when I say-- notice when I said expected payments, I said equilibrium-expected payments because, just from the auction format alone, we have no idea how much people are going to pay until we specify how people actually bid. So first, we need to choose or focus our attention on an equilibrium of the auction. And then, given that equilibrium, we can compute how much people are expected to pay.

So we're going to consider an equilibrium, a Bayes-Nash equilibrium of some generalized auction. And let's just make sure we understand what kind of object this is. So a Bayes-Nash equilibrium is going to specify functions-- let's see. Usually, we use star, but I think I'm actually going to use hats today because we're going to have to take some derivatives, and with the stars, it's a bit messy.

So bidder 1's equilibrium strategy is going to be a mapping from v lower bar to v upper bar to B1, all the way down to Bn hat. So remember, Bayes-Nash equilibrium is a strategy profile, and a strategy profile just specifies a strategy or a bidding function for each of the players. Notice here that, in general, these bid sets could be different, and our strategies have to reflect that.

So what does bidder 1 do? They know what their true value is. It's somewhere between v lower bar and v upper bar.

And to each true value, their bidding function specifies how much they bid. Well, what is a bid? A bid is just something in the set B1. Remember this is really abstract. B1 could be, I like it, or I don't like it.

And then my strategy would say, for which values, do I tell the auctioneer, I like it? And for which values, do I tell the auctioneer, I don't like it? And then, all the way down to bidder n, whose bidding function takes each value they could possibly have and assigns it to some bid that's feasible, some bid that's in the bid set that the auctioneer has presented.

OK, so we're considering a Bayes-Nash equilibrium, and what we're interested in is the payments. We're interested in the expected payment that each bidder makes. And we might also be interested in how likely each bidder is to win the good. So given this, we want to define what we might call the equilibrium-expected allocation and payments.

So let's first define the allocation. Maybe I'll call it Qi hat. Let's focus on bidder 1. I think it just makes the notation a little easier. Q1 hat of v1.

So the hat reminds us that this depends on the equilibrium v1 hat. And what I want to compute here is-- let's say we're playing this generalized auction, and we're following this Bayes-Nash equilibrium. Suppose I'm bidder 1. Suppose my true valuation for the good is v1. What is my probability or my expected probability of winning the item?

Well, whether I win the item depends on what other people's valuations are. So I need to take an expectation. So I'm going to take the expectation of-- well, let's write this.

Let's make sure we understand what's going on here. And I'll say this expectation is over v2 to vn. Why do I not take expectations over v1?

Well, I know v1. I'm bidder 1. I know what V1 is.

So this is what's, I think, really tricky about this notation. When you see v1 in here, that's just a number. That's something I know.

When you see vn, that's something I don't know, and that's something that's random. So think of v1. It's just something like 2. It's a number. We're not taking an expectation over it, but everyone else's bid-- everyone else's value, we are taking an expectation over it because we don't know what it is.

So what this is saying is if my value is v1 as bidder 1, well, then the amount I'm going to bid in equilibrium is B1 hat of v1. That's what my equilibrium bidding strategy tells me I should bid. Now let's take the example of bidder n.

Well, I don't know what bidder n's value is, but if their value is vn, then how much do they bid? They bid vn hat of vn. And we'll do the same for bidder n minus 1 and bidder n minus 3, all the way down to bidder 2.

So what's in here is the profile of bids that are submitted when we're following the equilibrium strategy. And bidder 1's value is v1, bidder 2's value is V2 all the way up to bidder n's value is vn. Now, bidder 1 knows her own value, v1, but she doesn't know these other values. So that's why we take an expectation. And then notice we have to feed in this profile of bids into our allocation function.

And I realized I made a mistake here. We need a 1 here. Should be q1 because this is the profile of bids. q1 of this profile is going to specify the probability that player 1-- or bidder 1 gets the item when this is the submitted profile of bids.

And then we compute expectations here. But I think this is a little abstract. So clear-- yeah.

AUDIENCE: So player 1 can't know her own allocation function, right? Because it depends on other people's values.

IAN BALL: So let's be clear. There's a key difference here between the function and the value of the function. So she does know the allocation rule. She knows if these are the bids that are submitted, then this is the amount I'm going to get.

But you're right, she doesn't know the bids that are actually submitted because she doesn't know how other people bid. So this thing in here, she doesn't know what value this takes. The value this takes is going to depend on these bids here.

But she's able to form beliefs over that. She's able to say things like, well, I don't know-- let's say there's two bidders-- I don't know what bidder 2's going to bid, but if bidder 2 bids 5, this is the probability I get the good. If bidder 2 bids 7, this is the probability that I get the good.

And because she has a belief about the distribution from which bidder 2's valuations are drawn and she believes that bidder 2's following bidder 2's equilibrium strategy, she's able to compute an expectation here. So she doesn't know the actual-- I guess what's also tricky here, it's weird to talk about-- this is a probability on the inside. And then we're also taking an expected probability. I know that's a little weird, but think of this number as basically always being 0, 1 on the inside.

So we're basically saying, what's the probability that I get the good? So we're taking an expectation of something that's sometimes 1-- it's 1 if I win-- and sometimes 0-- it's 0 if I lose. This tie-breaking stuff is not really going to come up. Any other questions? Great, great question, OK.

Great, so this is the equilibrium probability that I get the good if I'm bidder 1 and my valuation is v1. Now let's do the same with the transfer. So we wanted to find T1 hat of v1.

So this is exactly the object that we are computing over here. We are saying if I'm bidder v1-- if I'm bidder 1 and my valuation is v1, what is the expected amount that I'm going to pay in this particular auction, in this particular equilibrium of the auction. So again, we're going to take expectations.

Again, our expectations are going to be over v2 through vn because those are the values that bidder 1 does not know. And now we're going to write T1 of b1 hat v1 all the way up to bn hat vn here. And maybe I'll make it clear that this is known where these are uncertain from the perspective of bidder 1.

So we have bidder 1's bid all the way up to bidder n's bid. That forms the profile of bids. The auctioneer says, if I get this profile of bids, this is how much I'm going to ask player 1 to pay. And then we're taking expectations.

So we've now kind of formalized our question. Revenue equivalence tells us that this function, T1 hat of v1, seems to take the same form in a lot of different auction formats and a lot of different equilibria. In both the equilibrium of the first-price auction and the equilibrium of the second-price auction, this function was the same. And we'd like to understand why that is. So that's at a high level what we want to.

OK, well, we're going to have to use the fact that this is a Bayes-Nash equilibrium. So the fact that this is Bayes-Nash equilibrium tells me that-- let's take the perspective of bidder 1 who has a valuation of v1?

What we know is that bidding b1 hat of v1 is optimal. Given the way everyone else is bidding. I could bid whatever I want, but bidding b1 hat of v1 is the best possible bid for me. It maximizes my expected utility.

And a weird way you want to think about this is that tells me that bidding b1 hat of v1 is better than any other bid I could make. But these other bids, as you said, they're numbers. They could be words. We don't really know what these things are. But in particular, what this tells us is that bidding-- and this is the key observation-- bidding b1 hat of v1 is weakly better than bidding b1 hat of v1 prime for all v1 prime.

So in a sense, I haven't done anything here, but this actually turns out to be the crucial observation. So the point is, bidding b1 hat of v1 is better than any other bid. So in particular, what's another bid I could make?

I could choose the bid that I'm supposed to make in this equilibrium if my two valuation is v1 prime. v1 prime is just a different valuation. I could do that. But following my equilibrium bidding strategy must be better than this alternative bid b1 of v prime.

And then the name of the game here is just to try to write out mathematically what this means in terms of these functions here. So it's very useful to contemplate this particular alternative bid because we can use this alternative bid. We can express my utility from this alternative bid in terms of these functions here. And that's really the key step. So let's come back over here.

So I really want to emphasize everything we're going to do now is just following from this simple observation. So if there's any questions on this, I think this is key. Yes, thank you.

AUDIENCE: So does that imply that for every item in the bid set, there's some valuation that maps to it?

IAN BALL: That's a very subtle question. It does not imply that. So I think you're basically asking, are the bidding functions onto? Are they subjective functions?

If they were onto, then these statements would be equivalent. In general, the bidding function's not onto. So it's still the case that this implies this but not the converse.

So we're not using that assumption. If that assumption is violated, then we're actually throwing away information because we're saying, we know that this bit is better than every other bid, even bids that no one ever chooses in equilibrium. But in particular, we know this, so we're throwing away information potentially.

We're not even considering-- it's true that this bid is better than a bid that's never chosen, but we don't even need to use that fact. So we're not requiring that. Yes.

AUDIENCE: So why can b1 hat of v1 prime technically just represent all other possible bidding strategies?

IAN BALL: You said-- could you maybe-- I'm not sure if I quite follow.

AUDIENCE: Because there's infinite other bids.

IAN BALL: You're right. So I think the key here is I'm just using this forward arrow. I'm not saying that this-- that if this is true-- and this relates to this question-- I'm not saying if this is true, then this must be in equilibrium.

I'm only going in one direction. I'm saying, we know that this is an equilibrium, and an implication of the fact that it's in equilibrium, because it's equilibrium, I'm better off bidding b hat of v1 than making any alternative bid. So in particular, one implication of that is I don't want to bid 5. I don't a bid 3. I don't want to bid this.

Maybe we're not considering all the other possible bids. That's fine. Maybe we're throwing away information here, but it's certainly an implication of equilibrium that this must hold.

AUDIENCE: When you say throwing away information, what are you referring to?

IAN BALL: What I mean is the definition of equilibrium tells a lot of things, and I'm only using some logical implications of the definition of equilibrium. So it's like I have 10 equations, and I'm only using nine of them. And the 10th equation's true, but I'm just not focusing on it. I'm not using that fact. And maybe this all abstract stuff makes more sense once we go through the argument.

AUDIENCE: So you're saying like bidding b1 hat v1 is optimal. And there's other strategies that we're not comparing to. We're just comparing two. The bidding of b1 hat but with a different market.

IAN BALL: Exactly, so the fact that bidding b1 hat of v1 is optimal tells me a lot. It tells me that bidding b1 hat of v1 is better than any other bid I could come up with. But it's certainly better than these special bids, and I'm only going to use that fact here. Even though I know it's better than some other bits too, I'm not going to worry-- I don't even need to use that fact.

OK, so now, as we said, not equivalent but an implication of equilibrium is this observation. And now we just want to write this observation down mathematically using our notation here. So what we want to say is, what is my expected utility if I bid this? What is my expected utility if I bid this? And then we want to have an inequality between them.

So let's take the first. So what is my utility from bidding b1 hat of v1? Well, it's going to be v1 times the probability that I win minus my expected payment.

This is what my utility is. When I enter an auction, I say, well, this is my probability of winning. Whenever I win, I get my valuation for the good, but then I have to pay. Well, we'll just subtract my expected payment because we're assuming here people are risk neutral. So what is my probability of winning if I bid b1 hat of v1 using our notation over there?

AUDIENCE: Just a few hat of 1.

IAN BALL: Of v1, right? So this is q1 hat of v1. And then what is my expected payment?

AUDIENCE: It would--

IAN BALL: Right. Now the tricky part, what if I instead bid-- I deviate, and instead of bidding what I'm supposed to bid in equilibrium, I make a different bid. Specifically, I bid v1 hat of v1 prime.

Well, the crucial point is, what is my utility? Well, first, I still have a v1 here. Let's not get confused. My value is still v1.

The fact that I'm bidding differently doesn't change my value to v1 prime. My value is still v1, so it's going to be my value here, times, again, my probability of winning minus my expected payment. What is the probability of winning now going to be?

AUDIENCE: b1 hat v1 prime.

IAN BALL: Right, so I just plug in-- if I look carefully at this formula, if I replace b1 hat of v1 with b1 hat of v1 prime, well, then that's just the definition of this function evaluated at v1 prime. So I'm going to get Q1 hat v1 prime plus T1 hat. And this makes sense.

What's the definition of this? This tells me if I use bidding-- if I use the bid b1 hat of v1 prime, what is my probability of winning? And the crucial thing going on here is that my probability of winning only depends on how much I bid. It doesn't depend on my true valuation. It just depends on what I bid.

So the v1 prime here is really reflecting my bid, not my true valuation. So now what do we know? Well, we know this is greater than or equal to this. And maybe I've interpreted things so much, it's nicer to start a clean board and just write down our inequality.

So what do we know?

So what we know is that for all v1, v1 prime-- remember, the interpretation here is v1 is my true value. v1 prime is the value that I pretend to have when I bid. Well, we must have v1 Q1 hat of v1 T1 hat of v1 greater than or equal to-- so I've just brought that equation down here.

And maybe it's easier if we say for all v1, this holds for all v1 prime, just a slightly different perspective. Well, what does this tell me? Well, I have some function on the right side that depends on v1 prime, but it's always smaller than this left-hand side.

So in particular, what that tells me is that the value of v1 prime that maximizes this right-hand side must be V1 itself. So that tells me that V1, let's use our arg max notation is an element of the arg max over v1 prime of v1.

In words, what is this saying? It's saying if my true value is v1, I could bid as if my valuation is v1 prime for any other value of v1 prime. But because we know that I'm playing an equilibrium, the best thing for me to do is to bid as if I am myself, to bid as if my value v1 prime is exactly v1, and that's what this is telling us here.

And now let's just compute the first-order condition. So if v1 is a maximizer of this, then we should be able to take the derivative of this with respect to v1 prime, and then plug in v1 prime equals v1, and we should get 0. So let's do that.

So it's going to be a little tricky with the primes. You have to interpret the primes carefully. So this is v1 Q1 hat prime v1 prime minus T1 hat prime v1 prime. Evaluated at v1 prime equals v1, this should be 0.

Sorry, the primes, this prime on the inside is just a symbol. And this prime on the outside is a derivative. Sorry, that can be a bit confusing here. So you could call v1 prime z if you wanted. It wouldn't make a difference.

And what do we get? We get v1 Q1 prime v1. Minus, sorry. So this equals 0, or I could say it implies that they're equal.

Well, so a lot of work. But now we've narrowed things down a lot. I have a formula for the derivative of this expected payment that I make in terms of my probability of getting the good.

So now the trick is we're just going to integrate this. So if this equals 0, that tells me that this derivative equals this. So now we can just integrate this from v lower bar and see what happens.

Actually, maybe I'll just do it right here. So what does this tell me? It tells me that T1 hat of V1 is equal to T1 hat of v lower bar plus an integral from v lower bar to v1 of x-- I'll change to x.

Again, when we start integrating, you have to be careful. I've used a new dummy variable here, x. And there's a prime missing.

OK, so this is going to be the key revenue formula. Now, why is this so powerful? I think we want to step back.

And at first, you might say, OK, we just did a sea of algebra. What's going on? Well, what is this giving me? It's giving me a formula for, if I'm player 1 and my valuation is v1, and we're considering an equilibrium of this generalized auction.

It's giving me a formula for the expected payment that I make in this auction. But let's understand, in general, you might think that my expected payment would depend on a lot of things. You might think it would depend on the details of the auction and all these complicated things.

But if we look here, what is it depending on? It's depending on two things. Can someone interpret these two terms? These are the only things that appear on the right-hand side. So what is this first term?

AUDIENCE: The base cost of beginning the auction.

IAN BALL: Yeah, so it's just the expected payment that the lowest possible valuation person pays. That's the one thing, and I think you could interpret it maybe that way. And then what is Q1 prime-- I mean, the prime, let's just focus on the function Q1 hat.

What is this function Q1 hat? Can we interpret-- what does it mean? Yeah, maybe, I think. Yeah.

AUDIENCE: It's the marginal likelihood of winning for each infinitesimal change in your valuation.

IAN BALL: I think the integral is making you think about infinitesimal stuff, but really, it's-- oh, you're taking the derivative, exactly. Sorry, yes. So sorry, you said, yes, that's the derivative of it, exactly.

And if we take the derivative, it's just my probability of winning, exactly. And this is the infinitesimal, sorry, yes. Exactly, I forgot, yeah.

OK, so what it tells us is that a lot of the details of the auction don't matter. I mean, whether we're bidding what we name our bids-- notice the sets B1 through Bn don't even appear here. All that matters is the probability that I win the good in equilibrium and the amount I pay if I'm the lowest-value person.

So we immediately see that only a few details of an auction equilibrium pin down or uniquely determine the expected payment that everyone makes. So maybe I'll come over here.

So what's the implication? That the expected payment by, say, bidder 1 when the valuation is v1 depends only on-- well, T1 hat v lower bar, which is the expected payment of the lowest valuation-- more precisely, the expected payment that bidder 1 makes when her valuation is the lowest possible valuation.

And we had the derivative here, but if it depends on the derivative of Q1 hat, well, it also depends on Q1 hat. So it's a bit cleaner to say just depends on Q1 hat-- this function Q1 hat, which is the win probability. And this is an entire function. So it's the win probability of every valuation.

But the point is, in a lot of different equilibria of different auction formats, these two values take very-- these take very special values. So in many equilibria, in many auctions that we've looked at so far, what is-- so I'll just say often. Well, intuitively, what do you think T1 hat of v1 lower bar is often going to be?

AUDIENCE: 0.

IAN BALL: 0 because, in most auctions, if you have the lowest possible valuation for the item, in most equilibria, well, you're never going to win. And if you're never going to win, you don't want to pay anything. So often, auction equilibria have this property.

And now maybe this is a more subtle one. What about Q1 hat? Maybe I'll write Q1 hat of v1? Well, in a lot of auction equilibria we've looked at, which bidder wins the item?

It's the highest bidder. But another way of saying it is it's the bidder who bids the most, that's true. But it's also the bidder whose valuation is the highest. So often, when auctions work well, we might think that an outcome of the auction is that the bidder who values the item most gets the item.

But if the bidder who values the item most gets the item, then what is Q hat 1 of v1 going to be? Well-- Yeah.

AUDIENCE: Or as an expectation-- it will either be 0 or 1. So it's a kind of average of those probabilities.

IAN BALL: Its an expectation between 0 and 1, so it's going to be a probability between 0 and 1. Yeah, that's true. Can we say more about it? Yeah.

AUDIENCE: Maybe, but if there are n other players, then would be v1 to the power of n minus 1.

IAN BALL: In the uniform example, that's exactly what it would be. But more generally, maybe let's just write it would be the probability that everyone else's value is below v1 because if it goes to the bidder with the highest value, what's the probability that bidder 1 gets it when the valuation is v1? It's the probability that everyone else is value is below v1.

So it's the probability that the max of j not equal to 1 vj is less than or equal to v1. And we actually have notation for this. In the symmetric case, what is the notation? Yeah.

AUDIENCE: Big G.

IAN BALL: Big G, right? So this is also-- I'll write up here, big G of v1. So we often see that this holds. "Often" is in quotes, but what we see is that these tend to happen.

And if that happens, let's plug into a formula over here. So under these conditions, what does T1 hat of v1 give us? Well, what is this first term under those conditions?

AUDIENCE: 0.

IAN BALL: 0, OK, so we can forget that. Q1 hat 1 of x is what?

AUDIENCE: Small g of x.

IAN BALL: Right, so Q1 hat is big G. And then, when we take the derivative, we get small G. So what we get is the integral from v lower bar to v1 of x g of x dx. Does this look familiar?

This is exactly the formula we had. So now we understand where it comes from. So this is not a special property of first-price auctions or second-price auctions. Now we understand exactly when this is going to hold.

This is going to hold in any equilibrium of any auction that satisfies these two properties. Whenever my value is the lowest possible value, I pay nothing. That's the condition here.

And if the good is always allocated to the bidder with the highest valuation-- now, how that happens is going to depend on the details of the equilibrium, what the bids are. But as long as the auction has that property, then the expected payments are going to be the same. They're exactly going to take this form.

So indeed, it wasn't a coincidence. It wasn't just something about the first- and second-price auction. This is a very general fact. So I think it's good to understand-- I think sometimes when-- I should say exactly how it's done in the notes-- revenue equivalents can be stated in two different ways.

Sometimes people say only revenue equivalence only holds under certain conditions. Let's understand. This formula always holds-- or always. This formula holds quite generally, but it reduces to this special form under the additional conditions.

So you can still apply this formula to an auction where the good isn't always allocated to the person with the highest valuation or where the lowest-valuation person pays a positive amount. This formula still holds, but it reduces and simplifies under those further conditions. Yes.

AUDIENCE: Could you revisit the step of how we went through taking the derivative of our v1 equation?

IAN BALL: Yes.

AUDIENCE: --and getting to the revenue format?

IAN BALL: Yes, so from here to there? OK, so you're good up to this point? OK, so we're at this point. And it's a little spread out, but this is 0 over here.

So the 0 over here tells me that v1 times q1 hat prime of v1 equals T1 hat prime of v1. And now, yeah, I think the way I changed variables is tricky. I think the best thing to do would be wherever you see a v1, let's just put an x.

So this tells me x q1 prime of x equals T1 prime of x. Now I'm going to integrate both sides from x equals v lower bar to x equals v1. So when I integrate the derivative of T1 from v lower bar to v1, I just get t hat of v1 minus t hat of v lower bar. So you could think of that coming to the left-hand side. And then, when I integrate this, well, I exactly get this.

AUDIENCE: OK, that makes sense.

IAN BALL: Clear? OK. Thanks for asking. Any other questions?

So yeah, it was tricky. I used v1 here, but then v1 becomes my dummy variable x. And then I integrate up to v1.

Any other questions? So you might say, OK, this is a nice theory, but how can we ever use this? Well, it turns out this can be a convenient way to solve for equilibria in auctions. So let's go over a mini-- an application of revenue equivalence.

Maybe I'll call it RET, an application of the Revenue Equivalence Theorem. Let's look at all-pay auction. So let's do our simple case where we have just two bidders, v1 and v2. And they're in 0, 1, and they're uniform.

And now, instead of considering the first-price auction or the second-price auction, let's consider what's called an All-Pay Auction, APA. So in an all-pay auction, the allocation rule is the same as the first- and second-price auction. We allocate the good to the highest bidder.

So we allocate to highest bidder. But everyone has to pay their bids, whether they win or lose.

So I think it should be quite intuitive here that you don't want to bid very much in this auction because you know you have to pay your bid whether or not you win. So intuitively, people will bid less. But we might guess-- so let's guess that there's a symmetric equilibrium of this auction. Or let's look for a symmetric equilibrium of this auction.

So a symmetric equilibrium. Remember, a symmetric equilibrium means both players are using the same bidding function. So a symmetric equilibrium means we're looking for a function beta from 0, 1 to maybe 0 to infinity.

So we're looking for a symmetric equilibrium where both players follow the bidding strategy beta that says, if my valuation is this, I'm going to bid beta of my valuation. So this is what we're trying to solve for. I'll go over here.

And we're going to guess a symmetric equilibrium, but with a few more properties. I think we're going to guess that there's an equilibrium, and also, beta of 0 equals 0, and beta is strictly increasing. These seem pretty reasonable guesses.

So we think, because the environment's symmetric, maybe people use symmetric bidding functions. It seems to make sense that if your value is higher, you should bid more. That makes sense. And why do we expect beta of 0 to be 0? Yeah.

AUDIENCE: You wouldn't want to bid anything positive if the item isn't worth anything.

IAN BALL: Especially because it's an all-pay auction. I mean, that's always true, but in all-pay auction, if you bid something positive, what happens? You have to pay.

So you're paying a positive amount. And even if you get the good, you don't value it. So these seem like pretty safe guesses.

But let's say there is an equilibrium with this property. Will our revenue equivalent assumptions be satisfied? So let's check if there's an equilibrium of this property, what is the expected payment by someone whose valuation is 0? What's it going to be? Just 0.

If there's an equilibrium, where, if my value is 0, I bid nothing, well, then I'm going to pay nothing. It's going to be 0. And then if we're both using the same strictly increasing bidding function, well, which bidder is the good going to be allocated to? Yeah.

AUDIENCE: Whoever has a greater amount here.

IAN BALL: It's going to go to whichever bidder has the highest value because it's going to go to whichever bidder bids the most, but we're both using the same increasing bidding strategy. So who's going to bid the most? It's going to be the person who has the higher value. So it's going to be allocated--

So the two conditions we needed to get the revenue formula to simplify are satisfied. And that means we immediately know what the expected payment of each bidder has to be. So what does that tell us?

Well, what is-- let's just focus on bidder 1. It'll be the same for everyone else. What is T1 hat of v1 then?

Well, let's go through it. It's the integral. Well, what is v lower bar? It's just 0, 0 to v1 of x g of x dx. What is g of x here?

AUDIENCE: 1.

IAN BALL: It's just 1. Remember g equals f if there's two bidders. And if they're uniform, g is just 1.

So why? G is the density of the highest valuation among my opponents. There's only one of them. So it's just the density of the uniform distribution, which is 1. So what do we get?

Same formula that we always got before-- v1 squared over 2. Notice, this is exactly what we got way at the beginning of class. So in general, this is always the formula for the expected payment by a bidder with valuation v1 in our simple two-bidder uniform distribution case as long as our assumptions are satisfied.

But now I claim that we're done. What must the bidding function be? So let's step back and see what we've done.

We've guessed that this is the bidding function beta. We've said that if this is the bidding function beta, then the expected payment by bidder 1 when the valuation is v1 must exactly be v1 squared over 2. Yeah.

AUDIENCE: Is it 1/2?

IAN BALL: Even easier. It's an all-pay auction. So another way of writing the expected payment-- beta v1. If it's an all-pay auction, my expected payment is my bid because I always pay my bid. So of course, that's my expected payment.

If I bid 2, I always pay 2. If I bid 3, I always pay 3. That's my expected payment. So what's the answer?

AUDIENCE: j squared over 2.

IAN BALL: That's it, and we're done. So we've now found-- I mean, technically we should verify this. We've only said, this is a necessary condition, but it is true.

So you can check if it's right. Actually, it's v. We're done.

So revenue equivalence can be a convenient way of actually solving for an equilibrium because revenue equivalence gives us a very simple formula for the expected payment that each player makes. This is from revenue equivalence. The auction format gives us a relationship between the expected payment and the bid. In the all-pay auction, it's a very simple relationship. The expected payment I make is exactly my bid.

But now if we have two different formulas, two different expressions for T1 hat of v1, they must both be equal. And now we have this.

Maybe just one more example to see how we could derive the first-price auction equilibrium using this. Basically, we're going to do the reverse of what we did before at the beginning of class. Let's consider a first-price auction, again, in our simple example-- in our simple two-bidder uniform example.

Well, again, if we look for a symmetric equilibrium where if your value is 0, you bid nothing and beta is strictly increasing, then our two assumptions of revenue equivalence are going to hold. And therefore, we know-- so revenue equivalence tells us that T1 hat of v1 equals v1 squared over 2. That's still going to be true as long as the good goes to the bidder who values it most, and the bidder with 0 value pays nothing.

But now from the auction format, what is T1 hat of v1? Yeah.

AUDIENCE: That's the probability of winning times the--

IAN BALL: Exactly, so this should be my bid is beta of v1 times the probability that I win, which is v1. So now we set these equal. And indeed, we get-- which we already derived.

So we can go back and forth. We can solve for the equilibrium. We can see that revenue equivalence is satisfied. Or we can jump right to revenue equivalence and use that together with the rules of the auction to try to derive what the equilibrium bidding behavior will be.

And this is useful when we confront new auctions we haven't looked at before because of how general revenue equivalence is. I want to end with maybe one caveat about revenue equivalence, though.

And revenue equivalence is often misinterpreted. When people interpret revenue equivalence, they often say it's about the equivalence between auction formats. They say something like, the first price-auction generates the same revenue as the second-price auction.

But if we want to be really precise, revenue equivalence is a comparison between an equilibrium of one auction and an equilibrium of another auction. So a caveat, it's not about-- so it's not about the auction per se. It's about a particular equilibrium of the auction.

Maybe I'll say BNE of the auction. And in fact-- we'll give some examples-- in a given auction format, you could have multiple equilibria that generate different revenue for the auctioneer if those conditions aren't satisfied. So let's make sure we understand these conditions that the bidder with the lowest value pays nothing, and the good is allocated to the bidder with the highest value. These are not properties simply of the auction format.

They're properties of the equilibrium together with the auction format. If everyone used a different bidding strategy, then this may no longer be true. So these assumptions are not assumptions about the rules of the auction. They're assumptions about a particular equilibrium of the auction.

And whenever we apply revenue equivalence, we're not just comparing the revenue from two auctions. We're saying, let's look at equilibrium of one of this auction and equilibrium two of auction two, and let's compare the revenue there. So that's kind of a final caveat.

And we'll see some implications of that next week. So I'll stop there. Thank you.