[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: So today we're going to continue with our study of auctions. And we're going to study a bit of a richer model than the example that we looked at so far. But let's just recall what we've done so far with auctions.

So far, we looked at this very simple example where we just had two bidders, and each bidder's valuation for the item was uniformly distributed on the 0, 1 interval. So let's look at this example. We had our example where we had two bidders, and their valuations v1 and v2 were independent and uniform on the interval 0, 1.

And we looked at two different auction formats in this simple example. The first auction format would be a first-price auction. Remember, in a first-price auction, each bidder submits a bid. The person who submits a higher bid wins the item. And the amount that they pay is the amount of their bid.

And what we showed is that in a first-price auction, at least in this example, we had an equilibrium where each bidder bid exactly half of her valuation. So we had an equilibrium where maybe b1 star of v1 equals v1 over m and b2 star of v2 equals v2 over 2. So this was a symmetric equilibrium. Each bidder bids exactly the same way, given their valuation evaluation.

And of course, we noticed that each bidder is always going to bid strictly less than their true valuation. If they bid their true valuation, well, they can never come out ahead, because even if they won the auction, they'd end up paying exactly what they value the good for. So this way, there's a chance that when they win-- or when they do win, they're always paying strictly less than their valuation for the good.

Another auction you could run is a second-price auction. And in a second-price auction, the bidders submit their bids, just like in a first-price auction, and the rule for allocating the good is the same as in a first-price auction. The person who bids the most gets the item. The difference from a first-price auction is the payment rule.

So in a second-price auction, the winning bidder doesn't pay their own bid, but they pay the second-highest bid-- that is the highest bid among everyone else. And we argued informally that in the second-price auction, it's optimal to bid your true valuation.

So here, we have an equilibrium where b1 star of b1 equals v1, and b2 star of v2 equals v2. And I want to start by arguing why this is equilibrium, why you always want to bid exactly your true valuation in a second-price auction. So to see this, let's consider a Second-Price Auction, SPA, and let's take the perspective of bidder 1. So suppose I'm bidder 1, and suppose my valuation for the good is v1. And we want to look at different ways I could bid.

So we want to argue that it's always better to bid my true valuation v1 than to make a different bid. Well, how could I bid differently? Well, there's two rough ways I could bid differently. I could either bid more than my valuation or less than my valuation. So we're going to analyze those two cases separately.

So let's think about our bid space here. So this is 0. And let's fix our valuation here, v1. And we want to first see, well what if instead of bidding my true valuation v1, I chose an alternative bid, b1 prime that was, say, here.

So I'm contemplating a deviation where instead of bidding v1, I bid b1 prime. So I bid strictly higher than my true valuation. And we want to argue that this can never help me, that I'm never better off by bidding b1 prime than I am bidding v1. And to see this, well, what happens in this auction is going to depend on what the other bidder bids.

So I want to think of this axis here as representing what the other bidder is bidding. And then I want to compare the performance of this bid to this bid for each possible way that the other player bids. So let's first start down here.

Suppose the other player b2, the other player makes a bid b2 that is, down here. It's strictly smaller than v1. OK, so what happens if I bid v1? And what happens if I bid b1 prime? Let's start.

What if I bid v1 and the other player bids b2? What happens? What's the outcome of the second-price auction? Yeah, in the front?

AUDIENCE: You pay b2, right?

IAN BALL: You pay b2, and you win. Good. And you win the good. What if you had instead bid b1 prime? Yeah.

AUDIENCE: The same thing.

IAN BALL: Same outcome. So if the opponent's bid is down here, it doesn't make a difference whether I bid v1 or b1 prime. Either way, I win the item, and I pay the bid b2. So here, we'll say no difference between v1 and b1 prime.

What if b2 is up here? Now what happens if I bid v1, and if I bid b1 prime? Yeah?

AUDIENCE: You lose either way.

IAN BALL: You lose either way. So here, it also doesn't make a difference. They're outbidding me whether I bid b1 prime or v1, so I lose. I don't get the item, and I don't pay anything, so it makes no difference.

What if they bid in here? And we could worry about them bidding exactly V1 or exactly b1 prime, but that just adds some complications. So let's say it's strictly between v1 and b1 prime. Now what happens if I bid v1? I lose, because I've been outbid by my opponent, so I get 0.

If I instead bid b1 prime, well, this sounds pretty good, because now I win. I win the auction. But what happens if I win by bidding b1 prime when my opponent bids b2 here? Yeah?

AUDIENCE: You're paying more than your actual valuation.

IAN BALL: So I'm actually worse off here, because I could bid my true valuation v1, pay nothing, and win nothing, and my payoff is 0. Or if I bid b1 prime, I outbid my opponent. I win the item, but now I have to pay the second-highest bid, which in this case is b2. So I'm paying more for the item than it's actually worth to me. So I'm strictly worse off by bidding b1 prime than v1. So in this case, I'll say v1 is strictly preferred-- this is kind of a squiggly-- to b1 prime.

So there's three separate cases we've analyzed. In two of the cases, it doesn't matter whether I bid v1 or b1 prime. And in one of the cases I strictly prefer bidding my true valuation of v1 to bidding b1 prime. So the argument is almost done, but we only considered bids strictly higher than my valuation. We now have to go through the same argument and make sure that I never want to bid strictly less than my valuation. So let's go through the same thing over here.

So same figure. Here's v1, but now we're going to contemplate a downward bid. And you might think this is the more common bid when we played the-- or the more common mistake. When we played this first-price or the second-price auction, some people-- actually, most people bid pretty well, but some people bid strictly less than their valuation. So let's see whether this is a good idea. And let's again analyze our three cases. And again, I say three cases. We have to be a little careful if it's right here, but it's not going to change anything.

So case 1, what if my opponent's bid is down here? It's strictly lower than b1 prime and it's strictly lower than v1? Well, same thing. I lose either way-- sorry, I win either way. And I pay b2 either way, so there's no difference.

Up here, well, I'm being outbid. b2 is strictly higher than v1, and it's strictly higher than b1 prime. So whether I bid my true valuation or b1 prime, I lose the auction. I pay nothing, I get nothing. And therefore, it makes no difference. But again, the critical region is this middle region, where my opponent's bid b2 is strictly between b1 prime and v1. So what happens here if I bid truthfully and bid v1?

AUDIENCE: You win it and pay b2.

IAN BALL: Right. And what if I had instead bid b1 prime?

AUDIENCE: You'd lose.

IAN BALL: I'd lose. So I'm better off bidding truthfully, because I win and pay less than my valuation, so I make a strictly positive utility. Whereas if I had deviated and bid b1 prime, I would have lost the auction, won nothing, paid nothing, and therefore got a utility of 0. So once again, we see that v1 is strictly better than b1 prime.

So now we've completed our graphical argument to show that bidding truthfully is better than bidding anything else. And in fact, what we've shown is that bidding truthfully is actually a weakly dominant strategy. So normally, with auction equilibria in, say, first-price auctions, the best way for you to bid depends on the way your opponent bids.

In this case, bidding truthfully is best no matter how your opponent bids, no matter what strategy your opponent uses. Yes. Question?

AUDIENCE: I may have missed it, but--

IAN BALL: Yeah, fine.

AUDIENCE: --why isn't bidding b1 prime when your opponent bids b2 all the way at the bottom better than v1? Like don't you pay--

IAN BALL: Let's see. So there's six cases here. So you're in this one, case 1?

AUDIENCE: It's the graph on the right. And so you see--

IAN BALL: And here?

AUDIENCE: Yeah.

IAN BALL: OK. Great. Let's go over this. So let's suppose my opponent bids b2 here. What happens if I bid v1?

AUDIENCE: Oh, because it's second price.

IAN BALL: Because the second-price auction. It wouldn't be true if it was a first-price auction. Either way, I win. And the crucial point about a second-price auction is when I win, the amount I pay doesn't depend on what I bid. So the crucial property of a second-price auction is my bid can affect whether I win. But if I win, the amount I pay does not depend on my bid. It only depends on my opponent's bid. So in this case, if I bid here or here, I win either way, and I pay b2. Good. Great.

So now we've analyzed this example. I think you can see that this kind of argument is actually going to extend. We didn't really use the uniform distribution here. This is a much more general property of second-price auctions that it's optimal to bid truthfully.

But I want to make one kind of interesting observation here. We've looked at two different auction formats, and we've solved for an equilibrium in each auction format. So a natural question to ask is, which auction format is better for the auctioneer? Which auction does the auctioneer prefer?

Well, to say "prefer," we have to talk about the auctioneers preferences. Let's assume that the auctioneer likes money and they're risk neutral. So the goal of the auctioneer is to maximize their expected revenue in the auction. So now let's compute their expected revenue in the first-price auction.

Well, let's first just look at their revenue. Maybe we'll call it pi. Pi is a good notation for revenue. So what is their revenue in the first-price auction if bidder 1's value for the good is v1, and bidder two's value for the good is v2, and the bidders are following the first-price auction strategy. Well, that means bidder 1 is going to bid v1 over 2. Bidder 2 is going to bid v2 over 2. And how much revenue then am I going to make as the auctioneer? Yeah?

AUDIENCE: Whichever of those is larger.

IAN BALL: Right. So one way of writing that mathematically is just the maximum of these two. I'm going to make the maximum of v1 over 2 and v2 over 2. Why? Well, this is how much player 1 bids. This is how much player 2 bids.

In a first-price auction. I collect the higher bid because the higher bidder wins the item and pays their bid. So if v1 over 2 is bigger, they're going to-- player 1 wins the auction and pays v1 over 2. If bidder 2's bid is higher-- that is if v2 over 2 is higher, then bidder 2 wins the auction. They pay their bid v2 over 2, and this is what I get.

So this is telling me what my revenue is for a given realization of the valuations of the two bidders. But now I want to take expectations here. So my expected revenue is equal to the expectation of this. Maybe as a trick, I'll just pull out the half. I can always-- the max of v1 over 2 and v2 over 2 is just half of the maximum of v1 and v2. So maybe I'll write this as one half of the expectation of the max of v1, v2.

And this is an expression that comes up a lot for uniform random variables. It turns out the answer to this expectation is 2/3. But I want to give a little geometric argument for that. But let's just fill it in. If this is 2/3, and I multiply it by a half, then what I'm going to get is one third.

Where did this 2/3 magically come from? Well let's think. Well, first it's the number makes sense. I mean, v1 and v2 are both uniformly distributed on 0, 1. So this number should be between 0 and 1. That makes sense.

Now, the expectation of v1 by itself is a half. If you're uniformly distributed on 0, 1, the expectation is a half. So certainly, it makes sense that the expectation of the maximum of these two numbers should be higher than the expectation of just one of them. So we'd expect to get a number between one half and 1. So 2/3 makes sense.

Why do we get exactly 2/3? There's actually a really nice geometric argument that I want to show you. So we could think of v1 and v2 as living on this line. Maybe here is v2, maybe here is v1. And these are uniformly distributed.

The trick is actually to connect this line to make it a circle. So now let's view this as a circle, where this point is 0 and 1. So I've basically taken this, and I've just bent it around to look like this. So now I have v2 and v1 here.

Well, now, the maximum of v1 and v2 is going to be this length. But now, if I think about the circle, because things are uniformly distributed, I basically just have three uniformly distributed points on the circle. I mean, this one is a fixed point. It's not uniform, but it's as if it's uniform. It wouldn't make a difference.

So if I have three points on a circle, the expected distance between any consecutive ones is going to be one third. In expectation, this is a third, this is a third, and this is a third. So the expectation of two of the segments is going to be 2/3. It's maybe not quite a rigorous argument, but I think that should give you intuition for why this expectation is 2/3, because you just effectively have three points on the circle.

So now we see that in the first-price auction, the designer's expected revenue is one third. Let's now look at the second-price auction. So in the second-price auction, if bidder 1's valuation is v1, bidder 2's valuation is v2, and both bidders follow their equilibrium strategy of bidding truthfully, then what is the auctioneer's revenue going to be in this case? Yeah?

AUDIENCE: The minimum of v1.

IAN BALL: The minimum of v1 and v2. Because in this case, bidder 1 bids v1. Bidder 2 bids v2. The higher bidder wins the item, but they don't pay their own bid. They pay the second-highest bid, which in this case, what's the second-highest bid? It's the minimum of these two bids. If there's only two bids, the second-highest is the minimum of the two.

So here, we nicely see the trade-off. In a first-price auction, you get the higher of the two bids, but each of the bids is smaller, because the bidders anticipate that it's a first-price auction, so they shade their bids downward. In a second-price auction, you get the lower of the two bids. That hurts the auctioneer. But the benefit is people are bidding higher because they're bidding their true valuations. So the question is, which of these two forces wins?

Well let's compute this. We get the expectation of pi-- big pi, I guess-- of v1, v2 equals, well, just the expectation of the min of v1, v2. And if we use our little graph over there, can we guess what this expectation is? Yeah?

AUDIENCE: One third.

IAN BALL: One third, because now I have these three points on the circle. The minimum of the two is just going to be the length of the segment between two consecutive points. So it's going to be half as big as the length between-- over two segments. So I'm exactly going to get one third.

So what you see here is actually these two effects I described exactly offset. In the first-price auction, you get paid the highest bid, but people choose to bid less. In the second-price auction, you get the smaller bid, but people choose to bid more. And those two effects exactly offset. And the designer's expected revenue is exactly one third in each case.

Now, this was a very simple example with uniform distributions. So you might think, oh, this is just a coincidence. This is a fluke. It turns out this is a very deep and general phenomenon, that many different auction formats will generate the same expected revenue for the auctioneer. And that's something that we'll talk about a bit later in the course. Any questions on this? Yes?

AUDIENCE: The trick for finding the maximum and minimum, is that applicable if you had more?

IAN BALL: Exactly. Yeah. So maybe let's write this down. So if I did-- if everything was uniform 0, 1 and I had n-- I had v1 through vn, then the expectation of the max of v1 through vn-- and this is kind of a good fact to know-- guess, can you see what-- if we complete the pattern, we got 2/3 when n was 2. What do you think this is going to be now?

AUDIENCE: n over n plus 1.

IAN BALL: Exactly right. And in fact, you can actually go further. If I wanted the expectation of the second highest, it would be n minus 1 over n plus 1. The third highest, n minus 2 over n plus 1 all the way down to the expectation of the min equals 1 over n plus 1. So there's a very nice structure.

I could look at the lowest, the second lowest, the third lowest, all the way up to the highest. And those are exactly going to be equally spaced. And they're separated exactly by 1 over n plus 1. And you can use calculus to verify this, but I think this is a nice intuition for what goes on. And you could try to formalize that intuition. Good. Any other questions?

So now, as I said, this was a very simple example. You might worry that our conclusions are really sensitive to that. What we want to do now is study a more general auction setup, where instead of two bidders, we have n bidders, and instead of assuming their valuations are uniformly distributed, we allow for more general valuations. So we're going to look at a more general model.

So we're going to look at, in general, what's maybe called the symmetric independent private values model. Let me explain what these words mean as we go through. So now we're going to have n bidders. So we have bidders i equals 1 up to n.

And the bidders are symmetric in the sense that the distribution from which each bidder's valuation is drawn is the same. So of course, when we actually run the auction, someone might value the good more than someone else. But before we start the auction, we don't believe that the distribution of one player's valuation is higher than the distribution of another player's valuation. In reality, this is probably not a great assumption, but this simplifies things, and we can talk later about how it extends.

So what does that mean? It means that for each bidder i the valuation vi-- well, before we said the valuation was in 0, 1. Let's be a little more general. Let's say this lives in the interval between v lower bar and v upper bar. So this is the lowest possible valuation a bidder could have. And this is the highest possible valuation that a bidder could have.

Notice I don't have subscripts here because the range of valuations is the same for all the players. It doesn't depend on the identity of the bidder. And we'll say this is distributed according to a CDF-- that's a Cumulative Distribution Function-- capital F and a PDF, a Probability Density Function, little f. And I'll explain what these mean.

So this is just a way to describe a distribution. If you've taken a probability class, which is, I think, technically a prerequisite for this, you would have gone over this. But what are these things? Well, F of v just tells us the probability that the valuation is below v. So F of v is defined to be the probability that, say, vi is less than or equal to v.

So F is an increasing function. If I plug In v lower bar, I'm going to get 0, because the probability that vi is below v lower bar is 0. If I plug in v upper bar, then I'm going to get 1, because the probability that vi is less than or equal to v upper bar is 1. And then it's going to fill in between.

When I say it has a PDF f, What that means is I can express this probability as an integral from v lower bar to v of f of x dx. So that means we can visualize the distribution. Here's v lower bar, here's v upper bar by some density like this. And what this says, the peaks say these are places where the valuation is very likely to be. And then the valleys are places where the valuation is less likely to be.

And what we're saying is if we fix v here, the area under this curve is exactly the probability that the valuation lies between v lower bar and v. So what that means, one implication of that, is that it must be that the area under the entire curve is 1, because the probability that v lives somewhere is just 1. That's just a fact of probabilities.

So this is probably a little quick if you haven't seen this before. That is a prereq for the course, and I'm happy to chat after, or maybe you can self-study if you haven't seen probability distributions or densities before.

So symmetry refers to the fact that we're using the same distribution F and little f for all the bidders. It's not varying across the bidders. Independence means that v1 through vn are statistically independent.

This is the same assumption we made in the example. It's that if I know my valuation v1, that doesn't allow me to update my beliefs about other players' valuations. So what it rules out is things like positive or negative correlation, where when one bidder's valuation is high, everyone else's valuation tends to be high. Or when one bidder's valuation is low, everyone else's valuation tends to be low. This is also maybe a strong assumption, and we can, again, maybe talk later about how it could be relaxed. But this is what we're going to assume for today.

And I should maybe say one thing about these endpoints. I'm going to assume that 0 is less than or equal to v bar, strictly less than v upper bar. So the agents always assign a nonnegative value to the good. I think this is a pretty safe assumption. If they really hated the good, they could just throw it away. So we're always assuming that they value it at least 0. They could value it at 0, but they never value it at a negative number.

So we've talked about symmetry. We've talked about independence. What private values refers to-- and maybe I'll give us the acronym. This is SIPV. Private values refers to the fact that each bidder knows exactly how much he values the good. And bidder i's valuation doesn't depend on other players' valuations.

So you might imagine, if we were in an auction for a painting, and there was a chance that I could resell the painting, I might care about how much other people value the painting, because if they value the painting more, I might think, wow, I could resell it to them after the auction, and therefore, my valuation for the painting goes up. So we're ruling that out by assuming private values. We're assuming that I know my valuation for the good, and how much I'm willing to pay depends only on my signal about, say, how much I like the painting, and not on how much anyone else likes the painting.

So here we are, and we're going to try to analyze the first-price auction. We're going to try to solve for an equilibrium of the first-price auction. So goal, solve for BNE, a Bayes' Nash equilibrium, of the first-price auction in the SIPV setting.

I'll use a lot of acronyms-- of FPA, the First-Price Auction in the SIPV setting. Notice the setting and the auction are different. The setting is about who's involved and how much they value the good. The auction format, these are the rules that the designer gets to choose, or the auctioneer gets to choose.

The auctioneer can't choose how much people value the good. This is just a fact of the world. This is a fact of the setting. But the auction format is something that is chosen by the auctioneer.

Well, first, as always, we have to figure out, well, what is a strategy for each player? So what is a strategy for player i? Well, a strategy needs to say, how much does player i bid as a function of their valuation for the good. So a strategy for player i is going to be a function bi from-- maybe I'll write this.

So what does this mean? It means that for each valuation of the item that I could actually have, bi is my bidding function. Sometimes that's the terminology we'll use. My bidding function says, for each valuation, say, vi that I could have, my bidding function spits out bi of vi, which is the bid that I will make if my valuation is vi.

And just using the rules of the auction, we're assuming that bids can be any nonnegative number. I can't bid a negative amount. I can't ask the auctioneer to pay me. That would be a bit strange. But I can bid any nonnegative number that I want.

And an equilibrium is going to specify a strategy profile. So it's going to specify a bidding function for each of the n bidders. So as before, this is kind of a difficult problem because we're trying to solve for functions, not just single numbers. And these are really big. There's a lot of-- each player is choosing infinitely many numbers. How they bid if their valuation is 1, how they'd bid if the valuation is 2, 2.3, 2.5, there's too many things to solve for.

Before, recall that we got around this by guessing. We guessed that there would be an equilibrium where each bidding function was just a linear function. And that simplified things a lot. It reduced our problem to just solving for a single parameter. But that only worked in the special uniform example.

So beyond the uniform distribution, there's no reason to think that people are going to use a linear bidding function and bid exactly one third of their valuation, or one half of their valuation. So we have to be a bit more general. So what we're going to do is we're going to also make a guess. So we're going to guess, or look for, a symmetric increasing differentiable equilibrium. Let's go through what all these words mean.

Well, we're going to look for an equilibrium where every bidder uses the bidding function beta from v bar. so we're going to say, look, because the bidders are x antisymmetric, their valuations are drawn from the same distribution, we're going to look for an equilibrium where all the bidders use the same bidding function, the same strategy. So bidder 1, bidder 2, all the way up to bidder n, use this function beta. So if you wanted to be formal what we're saying is b1 equals b2 equals bn equals beta.

So everyone uses the same bidding function. And we're assuming that this bidding function is increasing. That makes sense. If you value the good more, it seems intuitive that you should bid more for the item because you care more-- you assign more utility to winning the items. You want to increase your probability of winning.

And this is just kind of a technical mathematical assumption. We want to be able to apply calculus. So we're going to assume that beta is differentiable. That is, we can always take the derivative of beta, beta prime, and that that derivative is always going to be strictly positive. That's basically increasing, a little stronger than increasing.

And now we've kind of reduced our problem. Instead of looking for n arbitrary bidding functions, we're now looking for a single bidding function beta. And we want to see if we can find an equilibrium where every single bidder uses this bidding function beta. Any questions here on our approach? Because we're going to get into a little math, so I want to make sure the approach is clear.

So what do we need to do? How are we going to solve for beta? Well, if this is going to be an equilibrium, it better be that if everyone else is falling the bidding strategy beta, the bidding function beta, it's optimal for me to also choose the bidding strategy beta.

So what we want to say is let's consider-- let's focus on bidder 1. It just makes it a little easier. So put yourself in the shoes of bidder 1. Suppose that-- I'm going to write-- define a function here U1 of b1 given v1 beta. Let me describe in words what this means.

I want to say I'm bidder 1. What do I know? Well, I know my true valuation for the good is v1. So this is my valuation. And let me suppose that all of my opponents are following the bidding strategy beta.

So what I'm supposing here is that b2 equals b3 equals bn equals beta. So I assume that player 2 is using the bidding function beta. Player 3 is using the bidding function beta, all the way up to player n is also using the bidding function beta.

And I want to say, OK, here I am. I believe that my opponent is using these bidding functions. I know my true valuation v1 for the good. What I want to compute is what is my utility if I were to bid b1. So b1 is just a number. This is not a function, it's just a number.

And once I compute this function, then I can start thinking about-- then I can maximize this function and find the optimal bid for me. But I need to put in this information because the best way for me to bid in this auction is going to depend on how much I value the good, and it's also going to depend on what I believe my opponents are using. So that's why I need all this notation here. And this is the key quantity.

So let's try to-- any guesses? Can we write out what this expression is going to be? Well, let's think it through. If I lose the auction, my payoff is 0, because if I lose, I don't get the good and I don't pay anything.

So we just have to figure out, well, what is my utility going to be? Let's suppose I win. If I win, and my valuation is v1, and my bid is V1, then what will my utility be if I win? Yeah?

AUDIENCE: V1 minus b1.

IAN BALL: Exactly, right. So this is my utility if I win. But I only get that utility if I do win. So now I have to multiply that by the probability that I actually win.

What is the probability that I win? Well, it's the probability that-- well, what does it mean for me to win? I win if my bid is higher than everyone else's bid. Now, there's always this question of ties. Because this bidding function is strictly increasing, any single bid by my opponents is going to have probability 0. So I don't really have to worry about ties. So let me just put in the probability maybe that my bid is strictly higher than all my opponents.

So that's the probability that, b1 is greater than-- well, it has to be greater than how much b2 bids, b3 bids, b4-- well that's just the maximum. Let me just take the maximum of all of their bids. So I'm going to say the maximum over maybe j-- let me write it out explicitly-- of beta v2 all the way up to beta vn.

I think, again, there's always this confusion. You see a probability, and then you see b1 and people think, b1 is random. No, I know what b1 is. b1 is just a number. What I don't know is how much my opponents are going to bid, because even if I believe they're using this strategy beta, I don't know what their actual valuations are. And their actual valuations affect how much they actually bid.

So if I bid b1, I'm going to win if player 2's bid is small, player 3's bid is small, player n's bid is small, all of their bids are small enough. That is, the maximum of all their bids is smaller than my bid. Yeah?

AUDIENCE: Because beta is strictly increasing?

IAN BALL: Here, I'm actually about to use that fact. I don't actually need to use it here. So I guess this statement doesn't use the fact that beta is increasing. Where that will come in is, when will beta of v2 be small? It will be small if v2 is small.

That's where increasingness comes in. If beta were not increasing, then maybe I'd win when my opponent's value was really high because that's when they bid low. So there will be a step where we use the fact that beta is increasing. But for now, we actually haven't used that fact. Because I win if my bid is higher than everyone else's bid, it doesn't matter whether their bidding functions are increasing or decreasing or crazy, but that will come up. Yeah? OK.

So now we will exactly go the next step. Let's analyze this probability here. Well, since beta is increasing, which of these players is going to bid the most? Well, beta is increasing, and they're all using the same bidding function. So which of my opponents is going to bid the most? Yeah?

AUDIENCE: The opponent that has the highest value.

IAN BALL: The opponent who has the highest value. So actually, I don't need to go all the way down. Let me just go here. I can simplify this expression as just writing this is beta of the maximum of v2 up to vn.

And here, I am using the fact that beta is increasing, because I'm saying the player who bids the most is exactly going to be the player who values the good the most. So in words, this is the maximum of the bids of my opponents. In words, this is the bid of my opponent who has the highest value. And because beta is increasing, these two are going to be the same. Great.

Now we can simplify it a little bit more. This is the random thing, and we want to focus on that. So let's apply beta inverse to this. So what we're going to get is v1 minus b1 times the probability that beta inverse of b1 is greater than.

So one way of thinking beta inverse is saying, OK, let's suppose I bid b1. What is the valuation that-- which valuation using bidding function beta would bid b1? And we're saying that valuation should be higher than everyone else's valuations. But mathematically, we're just taking an inverse

So now we've simplified things a bit. I think one key observation to point out is the math has gotten a little more complicated, but the basic trade-off hasn't changed. When I bid more, when I increase b1, the first term gets smaller, because when I win, I win less because I have to pay more. What happens to the second term if beta 1 goes up?

It's actually easier to see maybe here. If beta 1 goes up, this increases. Because if I bid more, it's more likely that I win. It's more likely that my bid is higher than everyone else's. So again, we face the basic trade-off, if I bid more, I increase the chance that I win, but I win less when I do win. And we have to balance those two effects.

So we're almost there, but now we need to analyze this expression. And what we need to do is we need to understand the distribution of the maximum of these random variables. So this is where we're going to use a little bit of math. We basically need to compute, what is the distribution function of this maximum?

So we'll define let's say G of v is equal to the probability that the maximum of v1 up to vn is less than or equal to v. This is a definition. And we want to try to understand this function G.

Well, what does it mean for the maximum of these random variables to be less than or equal to v? Well, that's the same as saying v1 is less than v-- ah, I shouldn't have v1 here. This should be v2. I apologize.

It's the same as saying v2 is less than v, v3 is less than v, all the way up to vn is less than v. For the maximum of things to be smaller than something, that's going to be true exactly if every single thing in the maximum is smaller than that thing. And the trick here is that we can use independence.

So it turns out this is the equal to the probability that v2 is less than or equal to v times all the way up to the probability that vn is less than or equal to v, because this is the-- this occurs if and only if each of these events occurs. And because they're independent, the probability of the joint occurrence of things is just the product of the probabilities. This is like saying the probability that I flip two heads is just one half times a half. It's one quarter. That's the fact that we're using if they're independent.

And again, I feel like whenever I go over this, if you've seen this before, this will be too slow. If you haven't seen it, it'll be too fast. I think there's not a good solution to that. Can I simplify this expression on the right side? Let's just be clear what my notation means.

I have n minus 1 numbers here, and I'm taking the product of them. This is one number. Then I have a number for v3, a number for v4, all the way up to vn. I have n minus 1 numbers and I'm multiplying them together. Yes?

AUDIENCE: So isn't that just the CDF at b to the n minus 1?

IAN BALL: Exactly, right? Because what is the probability that v2 is less than or equal to v? Well, that's just F of v. And the probability that v3 is less than or equal to v, that's just F of v. So I just get F of v multiplied by itself n minus 1 times because I'm going from 2 all the way up to n. So what I'm going to get is F of v to the n minus 1.

Let's just make sure this makes sense. Is G of v going to be bigger or smaller than F of v? And let's see if this makes sense. So F of v is a number between 0 and 1. So when we multiply a number between-- when we raise a number between 0 and 1 to the power n minus 1, that number is going to get smaller. And that makes sense because the probability for all of these things to be less than v is certainly-- that's less likely than the probability that a single one of them is less than or equal to v.

And we can also see another intuitive property of this. What happens when n gets really, really big? What happens to this right-hand side? Yeah?

AUDIENCE: Lower probability that you're on [INAUDIBLE]

IAN BALL: Yeah, so it gets really, really small, because the probability that everyone else is less than v goes down. So I just want to show this kind of makes intuitive sense. And then one final piece of notation that will be helpful is let me-- I'm out of space here. Let me go to a new board. Actually, realize it'll be good to be close to this board. So let's go to this board instead.

So I just wanted to find the density, the probability density function, associated with G. So let's define one other definition, little g of v is equal to the derivative of big G. Well, if big G equals this, now I just have to apply the chain rule, and I'll get n minus 1 times F of v to the n minus 2 times f of v. So I'm just using the power rule to bring the n minus 1 down, and then the chain rule to multiply by little f using the fact that little f is the derivative of big F. If you're a little rusty on this, that's not so important.

So with all these simplifications, I think we now have a pretty clean formula. We have U1 of b1 given v1 beta is equal to v1 minus b1 times-- well, we use our function G, And we get G of beta inverse of b1. So all I've done is used our notation G to apply to here. There's a little wrinkle because this was a strict inequality and big G is defined by a weak inequality. But the probability that the maximum takes a single value is 0, so that doesn't make a difference, so I can ignore that.

So now this is the expression we want to maximize, once again recognizing the trade-off that as I increase b1, this gets smaller and this gets bigger. So in words, this is the probability that I win. The way we're expressing it is the probability that I win is the probability that everyone else's valuation is below this magical threshold.

What's magical about this threshold? This is the valuation that they would have to have to bid exactly b1. So I'm going to win if every single player's valuation is below the valuation that they would need in order to bid b1.

So now what I want to do is I want to take the derivative of this expression so that I can choose the b1 optimally. I want to maximize this function. This function is telling me what my payoff is as a function of how much I bid. So let me just take derivatives and maximize this.

So if I take derivatives and evaluate the first-order condition, I'm going to get minus-- Let's go through what I've done. I've taken the derivative of the first term is negative 1 and multiplied it by the second term to get this. Then I've done the first term times the derivative of the second term. I've used the chain rule. So when I take the derivative of this, first I take the derivative of big G to get little g.

And then I need to take the derivative of beta inverse. But here I'm using what's called the implicit function theorem. The derivative of the inverse is the inverse of the derivative. It's the reciprocal of the derivative. So this is just a calculus fact. The slope of the inverse is just 1 over the slope of the original function. So I'm going to get this here.

So now I have a formula for what my optimal bid b1 should be. So my optimal bid b1 has to satisfy this formula. But I'm looking for an equilibrium. So it must be the case that this holds when b1 equals. beta of V1. This is kind of a subtle point.

Let's understand what we're doing here. This is a formula that's satisfied by my optimal bid. But in order for this to be an equilibrium, my optimal bid must be beta of v1, because we want it to be the case that if everyone else is using this bidding function beta, and my valuation is v1, then the optimal bid for me is exactly beta of v1. So we know that this must hold at beta v1.

So maybe a more formal way of saying this, so why is this true? Let me give the justification. It's because beta of v1 must be an element of the argmax over b1 prime.

So of all the possible bids I could choose as player 1, following my equilibrium strategy must be the best possible bid. So when I take derivatives of this and set them to 0 to find the maximum, that first-order condition must be satisfied by my maximizer, beta of v1. I think this is maybe a little subtle. This is kind of the key step. So any questions here? Yes?

AUDIENCE: So are you basically just saying that this is only necessarily true if b1 equals beta?

IAN BALL: Exactly right. Yeah. So I'm solving for the optimal-- this is a condition that must be satisfied by my optimal bid. But then I know that beta of v1 must be my optimal bid. So therefore, beta of v1 must satisfy this equation.

And that's going to simplify things a lot right. So let me plug in beta of v1 and see what I get here. And of course, the details of this are all in the notes, but I think if you don't see it in real time, it can get a little overwhelming.

So if we plug in beta of v1, we're actually going to get-- things are actually going to simplify a lot. We're going to get 0 equals negative G of v1, because beta inverse of beta of v1 is just v1. And then we get plus v1 minus beta of v1 times g of v1 times-- divided by beta prime of v1. So this is a much simpler formula. And we can actually simplify it a little further.

We know we're allowed to divide by this because beta prime is strictly positive. That's an assumption we made over here. So let's multiply by beta prime and move this to the other side. And what we get is beta prime of v1, G of v1 equals-- and this is a necessary condition for-- that our beta function must satisfy in order for us to have an equilibrium where everyone uses the beta function.

Now, I think we can actually intuit this pretty nicely. So here's some nice intuition for this. What I want to say is suppose my value is v1, and I'm contemplating, instead of bidding beta of v1, I could bid beta of v1 plus epsilon.

Well, let's think about epsilon being very small. So I'm supposed to bid beta of v1. We want to see what happens if I instead bid beta of v1 plus epsilon. And this is approximately going to be beta of v1 plus epsilon times b prime of v1. This is just a Taylor expansion.

So if epsilon is really small and positive, what I claim is that this is measuring the gain from winning more. And this is measuring the additional payment to understand what's going on.

If I deviate from beta of v1 to beta of v1 plus epsilon, then I'm going to win-- how is that going to change the probability that I win the auction? Well it's only going to make a difference if my opponents-- if the maximum value of my opponent is between v1 and v1 plus epsilon. If the maximum value of my opponents is between v1 and v1 plus epsilon, then when I bid as if my evaluation were v1, I lost. But when I increase my bid to bidding beta v1 plus epsilon, now I win.

The probability of that is roughly g of v1 times epsilon, because it's the density of my opponents' bid times the width of this interval. So what this right-hand term is exactly expressing is it's saying if I slightly increase my bid, I get my utility if I win times a slightly higher probability of winning, which is exactly g of v1 times epsilon. That's the amount by which my winning probability goes up by bidding a bit more.

On the other hand, if I bid a little bit more, I have to pay epsilon beta prime of v1 whenever I win. So I lose this much in additional payments whenever I win. And what's the probability that I win? Exactly big G of v1.

So there's a very-- if you could really had great intuition, you could maybe jump right away to this formula by thinking through this, but we derived it with calculus, so it doesn't matter. So to be very precise, what I'm saying is if instead of bidding beta of v1, I bid beta of v1 plus epsilon, then my expected additional payment would be approximately epsilon times the left-hand side, and my expected gain from winning more would be approximately epsilon times the right-hand side.

And if I'm already in an optimum, then the effect of this perturbation must be 0. This is like a first-order condition. And therefore, I get this equation. But that may have been a little subtle. So that's how I like to think about it. But if you don't see that the first time, I think that's OK. Yeah?

AUDIENCE: So the idea isn't that in the formula, you've-- so you have to multiply both sides by epsilon.

IAN BALL: Right. Yeah.

AUDIENCE: So epsilon is an implicit.

IAN BALL: Right. So the argument is if I multiply both sides by epsilon, then I guess the more formal argument would be I could actually choose epsilon to be a small positive number or a small negative number. So because I can choose it to be a small positive or negative number, the only way that I won't have a profitable deviation is if both sides of this are equal.

AUDIENCE: And then why does it make sense that some of this stuff is like derivative, some of it's like CDF, et cetera.

IAN BALL: Yeah, so let's understand the role of it. I'm saying, what happens if I bid slightly higher? Let's take epsilon positive. Well, the amount more by which I bid is epsilon times beta prime of v1. That's how much more I bid.

So you can think of that-- let's put the mini epsilon here. That's how much more I'm bidding. But what's the increase in how much I pay? Well, I don't actually pay that extra amount unless I win.

What's the probability that I win? It's the probability that everyone else's value is below mine. That's exactly G of V1. So this is the probability I win times the extra payment I make if I do win.

On the other hand, I'm looking at the extra probability that I win is g times epsilon. That's what I get down here. But then that marginal change in how much I win is multiplied by the utility I get when I do win. But I'll admit this is a little-- if you follow the math and you don't fully see all this intuition, I think that's OK the first time. Yeah.

So now we want to solve this. We're close. We have this formula. Let me erase the epsilons now. And now we have a differential equation. This is not really a course where we're trying to test how well people can solve differential equations, but this one, there's a pretty nice trick.

So let me just-- I want to get all the betas to one side because we now have a differential equation for beta. And now we want to find the function beta that satisfies this differential equation. So if we move all the betas to one side-- and let me also call my variable now v instead of v1, just to make things-- life a little easier. So then my differential equation becomes beta prime of v G of v plus beta of v g of v.

So this is my differential equation. So formally, what we've done is we've said-- we've guessed that there exists an equilibrium in which everyone used the bidding function beta. And then we've used the properties of equilibrium to argue that if that's the case, then we know a lot about what this beta function must be. In fact, it must satisfy this special differential equation.

And now we want to solve for this differential equation. Find the function that satisfies this. And if anyone's recently taken differential equations, they may see there's kind of something special about this left-hand side. Does anyone-- anything pop out at you. I don't know you just took a class on ODEs? Yeah?

AUDIENCE: That looks like the product [INAUDIBLE].

IAN BALL: Exactly. So let's express the left-hand side as beta of v G of v-- this is kind of sloppy notation, but I mean the derivative of this whole thing. Maybe I'll write d/dv. I don't know. Everyone has different notation for derivatives. Let's write that.

Well, now I have the derivative of something, and I have something else. So now I'm just going to integrate this from v lower bar to v. There's always an issue with what terms we use. But, OK, so let's integrate both sides from v lower bar to-- well, we don't want to call it v because we already have v here. So let's call it V hat

So on the left-hand side, we get beta of v G of v evaluated from v equals v lower bar to let's say V hat. And on the right-hand side, we get the integral from v lower bar to V hat of v g of v dv. So if these two sides are equal, certainly the integrals of the two sides are equal. I've integrated the right side from v lower bar to v upper bar. And I've also integrated the left side from v lower bar to v upper bar.

But this just takes a really simple form because I'm taking the integral of a derivative. The integral of a derivative is just the antiderivative at the endpoints. And this notation here means I plug in v equals V hat, and then I subtract the expression with v equals v lower bar. I think this is pretty standard notation. So maybe I'll put in arrows and I'll write integrate.

So now we're getting closer. But I'd argue that this left-hand side simplifies even more. How can we simplify this left-hand side? What happens at v lower bar? Yeah?

AUDIENCE: G of v lower bar is 0.

IAN BALL: G of v lower bar is zero. So we can just ignore that term. That disappears, right, because we get 0 times something. So we only get the top term. So we get beta of V hat times G of V hat on the left equals the integral from v lower bar to v upper bar of v g of v dv. And now we've solved our differential equation because we have beta by itself. And now we can just divide by big G. So let's divide.

So we get our solution is beta-- and this big V is annoying. So maybe I'll change notation a little bit and write beta of v-- let me just use v instead of V hat and then maybe x instead of v. So let me say this is the integral from v lower bar to v of x g of x. No. So let me check wherever I had a v hat I want a V. And wherever I had a little v I want an x. Looks good.

So here's our solution. So what we've solved for is the bidding function that every player is going to use in this symmetric equilibrium. I should warn that-- I should caution that there's one kind of trick we've done. Technically, what we've shown is that if there's an equilibrium of this form, then beta must look like this.

But really, we have to check that if beta looks like this, it's actually an equilibrium. We've shown a necessary condition, not a sufficient condition. But you can check that this actually is an equilibrium, so that's OK. It's sort of like when you solve a maximization problem and you set the derivative to 0, well, sometimes you don't have a maximum, you have a minimum or something. That could be going on here, but we can check that that's not what's going on.

So let's try to interpret this. Well, we already have a formula for the equilibrium in the special uniform case. So let's see if this agrees with our answer there. If it doesn't agree there, then we're kind of in trouble.

So in the uniform 0, 1 case, well, what is g of x, and what is G of v? Well, let's remember what g is. g is the distribution of the maximum value of my opponents. But if there's only two of us, then the maximum value of my opponents is just the value of my opponent. So in the uniform 0, 1, with two bidders, what is g of x going to be?

It's going to be the density of the distribution of my opponent's value. And what is that density in the uniform case? Just 1. And what is G of v going to be? Well, it's the CDF of the uniform distribution. So that's just v. The probability that a uniform distribution is below v is just v assuming v is between 0 and 1.

So let's plug this in. And I guess the final step is v lower bar is equal to 0. So these are all the special restrictions we get in this uniform example with two players-- I should say two bidders. Let's see. I guess we'll see in real time if we've made a mistake. I don't think we have, but we'll see.

So we get the integral from 0 to v of x over v dx. x times 1 divided by v dx. Well, that's just 1 over v is a constant. We can take that out of the integral from 0 to v of xdx. What is the antiderivative of x? It's just x squared over 2. So we get v squared over 2 times 1 over v.

And, phew, we didn't get it wrong. We got v over 2. So that's good. So this formula is consistent with our guess in the very special case. But this is a much more general formula.

Now, how can we interpret this formula? There's actually quite a nice way of looking at this. Does this fraction remind you of anything? I don't know. Does anyone see a way of interpreting this integral in terms of conditional expectations?

So G of v is the probability that all my opponent's valuations are below v. And g of x is the density of the distribution of the maximum of my opponent's valuations. So it turns out that this is actually a conditional expectation. And, again, this is a probability fact.

So it turns out-- what is this? It's actually the expectation of the max of v2 through vn conditional on the event that this maximum is less than or equal to v.

So this gives us a nicely interpretable formula. What does this say? How much do I bid in the equilibrium of the first-price auction when my value is v? What I'm actually bidding is the expectation of the highest value of my opponents, conditional on the event that the highest value of my opponents is less than or equal to v, less than or equal to my true valuation.

And there's a nice connection here. Maybe we'll try to do it. But let me pause, see I feel like I've lost people a little bit. So any questions? I know it's been a little mathier than we've normally done. So any questions here? Yeah, let's pause. Yeah?

AUDIENCE: So that the condition is that your valuation is higher than the valuation is like-- higher than the valuation of the maximum opponent. What if that's not true?

IAN BALL: So you're asking about this final step? So let's understand right here. v is just a number. So what I'm writing here is just a formula for just something that's defined as a function of this number v. So mathematically, all this is is a very special conditional expectation where this is a random variable. This is a random variable. And I'm just looking at a very special conditional expectation of a random variable.

So it may be that their valuations are lower than mine. What we're just saying is it's kind of a thought experiment. It's saying how much am I going to bid? Well, one way I can compute my equilibrium bid is by computing this very special conditional expectation.

But you're right that the true value of my opponents may be lower than v. This is just kind of a formula for what my bid will be. Yeah.

AUDIENCE: But it could also be that the true value of your bid is higher than v, right?

IAN BALL: It could be. Yeah.

AUDIENCE: But you're just saying like you're basing it off of the idea that if everything were less than or equal v, like [INAUDIBLE].

IAN BALL: I can just define this conditional expectation. So for now, I mean, it's just a definition. It's just a function of v, and it's a number we can compute. But it gives us kind of a nice interpretation that I'll go through in a second. Yeah. Yes. In the back.

AUDIENCE: Can you explain one more time why g of x over G of v is the--

IAN BALL: Yeah. OK. And this is something-- yeah, if you haven't seen before. So this is--

AUDIENCE: Wouldn't it also be the x

IAN BALL: Say it again.

AUDIENCE: Instead of just g of x, would it-- because x g of x is an expectation.

IAN BALL: So this is an unconditional expectation. When I put the g over v here it becomes a conditional expectation.

AUDIENCE: I wonder why x is [INAUDIBLE]

IAN BALL: Oh, I guess-- so it's a question of what we circle. So I'm saying the conditional expectation can be expressed as x times the conditional density. So what I've circled here is the conditional density. So this is the conditional density given the event that everyone else's value is below a certain thing.

So maybe I'll just say, what's the general property I'm using here? Let me go down here and write this kind of math fact. So if you have a random variable X distributed with a density g, then we can also compute the conditional density. So we can say what is the conditional density g of x given that X is less than or equal to v And that's defined to be g of x divided by the probability that X is less than or equal to v, which is exactly g of x times big G of v.

One way to think about this is if I want to understand the conditional probability, or the conditional distribution of x, given that x is below v, well, it's still proportional to the density, but now it has to integrate to 1 over this condition of n. So conditional on my opponent's value being below v, the conditional density, if I integrate it from v lower bar to v, better give me 1, because that's how conditional distributions work.

And if we do this integral from v lower bar to v of g of x, we exactly get big G of v. And then if we divide it by big G of v, we get 1. So this is just-- it's a normalization to ensure that the distribution, the conditional distribution, integrates to 1.

But again, as I said, we're using a little more probability. This is maybe the one lecture where we're kind of leaning on some of these official prereqs that we haven't really enforced. And yeah, some of these things may take a little refresher. But this is the general fact that I'm using. Any other questions?

So one final observation. If we try to look at this equation star-- let me go over here. So we said that if my valuation is v, this is how much I bid. So if I win the good, the amount that I'm going to pay is exactly beta of v, which is exactly this formula over here.

But let's just go through a little thought experiment. Let's suppose that we instead-- we did a second-price auction instead of a first-price auction. Well, now everyone's going to bid truthfully.

So let's say the profile of valuations is v1 through vn. If we're all bidding truthfully, and we're now in a second-price auction, when is bidder 1 going to win? What must be true about v1 through vn for bidder 1 to win the item? Yeah?

AUDIENCE: V1 is the max [INAUDIBLE]

IAN BALL: Yeah. So maybe another way I'll say it is-- so bidder 1 wins if V1 is greater than or equal to-- or maybe I'll say greater than-- I'll say greater than or equal to the max of v2 up to vn. And if they win, what do they pay?

How much does bidder 1 pay if they win the auction? Well, bidder 1 pays the second-highest bid, right? The second-highest bid is going to be the highest bid of his opponents. But his opponents are bidding truthfully, so the second-highest bid is exactly going to be the maximum of v2 up to vn.

So in a second-price auction as opposed to a first-price auction, the amount that I pay is random. I don't know how much I'm going to pay because it's going to depend on the valuations of my opponents. But the amount that I'm going to pay is exactly this. And I'm going to pay that exactly if I win.

So what is the expected amount that I'm going to pay given that I win? It's exactly going to be-- so suppose I win. What is my conditional expected payment? It's exactly the expectation of this, conditional on the event that this is less than or equal to v1.

Wait a second. That's exactly the formula we got over here. So this is an expression of revenue equivalence. It says that in a first-price auction and a second-price auction, the expected amount that each type of bidder pays is actually going to be the same. Here, I bid B of v. It's just one number, and I always pay that if I win.

Here, how much I pay depends on what other people bid. I don't know what I'm going to pay. If other people have valuations really close to mine, I'm going to pay a lot. If they all have valuations really low, I'm going to pay a little. But the expected amount that I'm going to pay is actually going to be exactly the same as the amount that I pay in the first-price auction.

And this is the first indication of this revenue equivalence fact, that while the way that we bid is going to be different, the expected amount that each bidder pays is going to be the same. And for that reason, the auctioneer's revenue, expected revenue, is also going to be the same. So this is the first hint at something called revenue equivalence, which we saw in a special example before, and now we see is a much more general phenomenon. But we'll talk about that a bit more next class.

I guess it's a holiday on Tuesday, so I guess I'll see everyone next Thursday. And I'm happy to stick around with any questions. Thanks.