[SQUEAKING] [RUSTLING] [CLICKING]

IAN BALL: So today we're going to continue with Bayes-Nash equilibrium. Last class was a bit of a theoretical class where we defined the concept, and we went over the foundations. Today, we're actually going to solve some examples. So we're going to solve an example of Cournot competition, and then we're going to get started with auctions, which is one of my favorite topics, and we'll cover that for a few lectures.

So let's start with an example of Cournot competition with unknown demand. So you recall that the model of Cournot competition we've introduced so far, it was assumed that both firms exactly knew the demand curve. They exactly knew what the price would be as a function of both of the quantities that was set. But that's kind of an unrealistic assumption.

So let's now suppose that the demand or, more precisely, the inverse demand, P of Q. Before we just said it was 1 minus q. Let me write what we had before, and then I'll change it. Before, we wrote something like this, so that if the quantity Q was produced on the market, then the price was either 1 minus Q, unless Q was really, really high, in which case, the price was 0.

But what we're going to look at today is what if this intercept that used to be a 1 is now a theta, and theta is unknown, or theta is going to be private information. So we're going to look at, I guess, the simplest possible case, where one of the firms knows the value of theta, but the other firm does not.

So let's assume that-- let's say theta can take two possible values. And let's assume they each occur with probability 1/2, so kind of the easiest case. So theta is an element of, well, either theta or theta high, and each with probability 1/2. Great.

So of course, we're going to assume here that the labels make sense. So theta L is smaller than theta H. So if the state is theta H, this is a high-demand state. And it means that for a given level of quantity that we put on the market, the market price is going to be higher. Because it turns out that more people like the good than we necessarily expected.

If the demand state is theta L, then that's the lower number here. And for a given quantity Q that's produced, the price is going to be lower because there's going to be less demand for the product.

So this is certainly relevant when a firm is introducing a new product. They don't know how popular it's going to be. They may have some uncertainty about that, and that's what this is capturing here. And we're going to assume that firm one observes theta, but firm two does not.

So one story could be that firm one is kind of an experienced incumbent. They know how fashions change, and they're able to predict what demand is going to be. Firm two might be this new entrant who doesn't have as much data, doesn't have as much experience, and they don't know what theta is.

Now, of course, what is assumed is that this distribution is common knowledge. So firm two knows, or believes, that the state is equal to theta L with probability 1/2 and theta H with probability 1/2. They know that, but they don't know exactly which value it is.

And moreover, everything I've written on the board is common knowledge. This is this subtle thing. So firm two knows that firm one knows the state. And we're going to now write this down for you.

So now let's try to solve for-- and I guess the last component I need to say, maybe I'll just say it here, is let's assume costs are 0 just to make it the algebra a bit easier. So marginal costs equals 0.

So as usual with our Cournot specification, your profit is your quantity times your profit per unit, which is price minus cost. In this case, cost is 0. So your payoff is just your quantity times the market price.

So now I think the first step, so now our goal is solving for B and E. A Bayes-Nash equilibrium-- I think sometimes the hardest part, often the hardest part of this exercise is figuring out what a strategy is. So the first step is to figure out what a strategy is. And once we've done that, I think we're in pretty good shape.

So remember, a strategy is going to need to specify an action a player is going to take as a function of whatever their mental state is. If they have two distinct mental states, they can take different actions in those states. So any ideas here? Let's look at firm one and firm two. What are the strategies going to be for each of these players? So not what the equilibrium is, but just what kind of object is it going to be? Yeah?

STUDENT: The strategy will be some non-negative yield number.

IAN BALL: So a quantity. And that will be a strategy for whom?

STUDENT: [INAUDIBLE] the strategy.

IAN BALL: So close. So not quite. So I think that's right for firm two. So firm two is going to be choosing a non-negative real number. You're right. We have to be careful here. So you're right.

Firm one will wake up. They'll see that demand is either high or low, and they will actually choose either-- they'll choose a non-negative real number. But their strategy has to be a complete contingent plan. Yeah, you want to follow up on that? Yeah?

STUDENT: There should be [INAUDIBLE].

IAN BALL: Exactly.

STUDENT: There should be two values in the [INAUDIBLE], one based on theta low and one-- so one for theta low, one for [INAUDIBLE].

IAN BALL: Great. So I'm glad you answered that way. Your first answer was a really common thing that we see in exams. And also, it's good that this came up.

It's tricky because firm one will actually, in reality, choose a single number, but their strategy has to specify two numbers because it must be a complete contingent plan. And if it didn't specify two numbers, there's no way firm two could figure out what to do because firm two doesn't know the state of the world. So if they don't know, or don't form beliefs about how firm one is behaving in each state of the world, they're going to be in trouble.

So for firm two, it is just a number. Maybe we'll call it Q2. Maybe I'll just say it's a non-negative number here. But firm one can, in principle, produce a different amount in situations or in states when demand is high and in states when demand is low. So they're going to choose two numbers, Q1 of theta L and Q1 of theta H. And these are both going to be non-negative quantity numbers.

If we want to think of this in terms of our formalism from last week, firm two has a single type. They have a single mental state. Sometimes we call that T bar. And when I just write Q2, that's really Q2 of T bar. It's the quantity they produce for their single mental state.

Firm one actually has two mental states. And the way we wrote it last time is we would have written it something like this. We would have said, firm one has type TL. And when their type is TL, they know the state is theta L. And when their type is TH, they know the state is theta H.

And if we wanted to be really formal, we would write Q1 of TL and Q1 of TH. But then you just have to keep track of the T's and the thetas, and we're going to take a shortcut and write it this way. But the understanding is that the reason we write theta L here this is the type of firm one who knows for certain that the state is theta L. I want to be clear about that interpretation.

So let me put these in parentheses. So now we want to solve for B and E. But I think first it's kind of good to always try to guess and form intuition about what our answer is going to be.

So in equilibrium, how do you think Q1 of theta L, Q1 of theta H, and Q2 are going to be ordered? Any guesses about how those three numbers might be ordered? See if we can form some intuition in advance.

So imagine you're a firm. If you think this product is very, very popular, and the market price is going to be very, very high, are you going to produce more of the product or less of the product? More of the product.

So let's start up here. Between Q1 theta L and Q1 theta H, which one do you think is going to be higher? We think this is going to be higher. And it will be. Now maybe a trickier question, how do these two compare to Q2? That's maybe a bit harder.

STUDENT: Somewhere in the middle.

IAN BALL: Somewhere in the middle. Yeah, exactly. So it turns out, and this is maybe not obvious and depends a bit more in the details the problem, but if I'm firm two, I know there's a half chance the state is theta L and a half chance the state is theta H. So it's intuitive that the amount I'm going to produce is somewhere in between.

Now, it's a bit trickier because when the state is theta L, it's also going to change what firm one produces. So that conclusion is a little sensitive to the specification. But in this case, that's going to be true. So let's try to write it down.

So we've gone through step one which is identifying the kinds of objects that we're choosing-- how many numbers we have to choose, how many functions we have to choose. The next step is to write down the equilibrium conditions. And the crucial thing about BNE, Bayes-Nash Equilibrium, is that we're going to have a condition, an optimality condition for every type of every player.

So remember, BNE, Bayes-Nash Equilibrium, requires optimality for every type of every player. And that means we're going to have a different equation, a different optimality condition for every type of every player that says that type is behaving optimally.

So in this problem, how many conditions are we going to have? Three, right? Because there's one type of player two. So we're going to have one condition for player two that says player two is behaving optimally. And then we're going to have two different conditions for firm one saying if firm one is the high type, and they observe that demand is high, then they're behaving optimally. And if they observe that demand is low, then they're also behaving optimally.

Now, to write out these conditions, it's nice to have a little math notation. People may have seen this, but I want to introduce it to make sure we're on the same page. So we often write things like max-- let's just go through a simple example to define this notation.

We might want to maximize some function f over x. And let's put some range in here. Let's say x is between negative 1 and 1. So what this expression denotes is the maximum value of this function over this range.

There's a different expression, a different notation that we want, which is the argmax-- let's use this-- of f of x, which is the set of x values, the set of inputs that achieve this maximum. And the easiest way to see it is just to go through an example.

So let's say f of x is x squared. I think just doing a numerical example is the easiest thing. So maybe I'll just erase the f of x to avoid confusion. Let's just-- I don't have too many equal signs.

So what is the maximum value of x squared over the domain where x ranges from negative 1 to 1? This is just 1, right? So this is not testing the math, but just the notation. This is the number 1, because that's the highest value that this-- actually, let me use two to avoid confusion. Sorry. Because otherwise, that would not be a good choice. So the answer is 2.

What about here, the argmax of 2x squared? Where is this maximum value of 2 achieved, at what x points?

STUDENT: 1 and minus 1.

IAN BALL: 1 and minus 1. So when you write argmax, this is actually a set consisting of the two numbers 1 and minus 1. So it's good to distinguish. This max says we take in a function, we take in a domain, and it gives us the highest value that function takes over the domain.

When we write argmax, now what we're returning is not the value, but an input to the function, an argument of the function. And we're saying at which arguments in the domain does the function achieve this maximum? So this is helpful notation because often we'll write things like x star is an element of the argmax.

And what this means is if you were someone who were choosing x, and you wanted to maximize this function, you could choose either negative 1 or 1, and those would both maximize the function. And the way we would write down that condition is to say x star, the thing that you choose, is an element of the argmax. It's an element of the set of things that maximize the function. So we'll use that notation when we write down our equilibrium conditions. Any questions on this?

Basically, the advantage of this is we only have to write the function once. If we had to say f of x is greater than or equal to f of x prime for all x prime, you just have to write things twice. And when the formulas get long, this is a nice way of writing it.

So now after that digression, let's now write out our optimality conditions for this game. And as we said, there's going to be three conditions. So we're going to have the condition for maybe, I'll say, F1-- this is firm one-- when they know the state is low.

There's F1 theta H-- firm one when they know the state is high. And then there's F2, firm two, and they don't know anything. There's only one type of them. So we just need a single optimality condition for the firm.

And we're looking for an equilibrium, so we're looking for values of this that form an equilibrium. So we'll probably put stars in here. So what we're going to get is this system of equations is going to depend on, maybe I'll call it Q1 star of theta L, Q1 star of theta H, and Q2 star. Those are the three things we're trying to solve for.

So let's start with theta L. So what is firm one choosing? They're choosing a non-negative quantity. So what they're choosing is Q1 star of theta L. And we want this to be in the argmax of something because we want this to be an optimal choice for this firm when they know that demand is low. So we're going to say this is in the argmax over Q1 prime.

So what I'm saying here is I'm-- what I want to fill in here is, what is the payoff to firm one? If they know the state is theta L, they choose quantity Q1 prime. And they believe that the other firm is following their equilibrium strategy.

So here, we're going to have the expected revenue or profit of firm one. And to say that Q1 star of theta L is in the argmax is to say this level of production maximizes the revenue of firm one. So what is it going to be?

Well, if I produce Q1 prime, then my quantity is Q1 prime. But then I have to multiply that by my revenue-- let's see-- let me just write it like this-- by my revenue per unit. So that means I need to plug in the market price. And we know that the market price is going to be theta minus Q.

So if I'm firm one, I know the state is theta L. So I'm going to put in theta L here. And then I need to subtract the amount that I produce minus Q1 prime. What else do I have to subtract? Yeah? Say it again?

STUDENT: Q2 star.

IAN BALL: Q2 star. Exactly And this is it. So let's understand because I know the state is theta L, I use theta L when I'm computing what the market price is going to be. If I'm firm one, and I choose quantity Q1 prime, and I believe that my opponent is following the equilibrium strategy of producing Q2 star, then this is going to be the market price. And I'm going to make that for every unit I produce, which is Q1 prime, and that's going to give me revenue.

Now we're going to get a similar formula for theta H. And I want to make it clear here-- sorry-- Q1 prime is a number here. So this might look like a function being evaluated somewhere. That's not what I mean. It's just a number times another number. And the primes don't mean derivatives. They're just notation.

So here, it's exactly the same. But now the market demand is different. So the key difference here is even if we both produce the same amount, the market price that we face is going to be different in the high state than it would have been in the low state.

And now from two, this is a bit of a harder one-- so they're going to say the quantity I'm actually producing is optimal among all the quantities I could produce. But now they're in a bit of a tricky situation because they don't know what the state is.

So firm two needs to take expectations over the state of the world. So there's a half probability that the state is high. And if the state is high, and firm two chooses quantity Q2 prime-- let's write that in here-- the question is, what is the market price going to be in the case that the state is high in this equilibrium?

Well, now we have, I guess, two-- the market-- the state matters in two ways. Because it's going to affect the actual intercept in the market price or the inverse demand curve, but it's also going to affect what firm one chooses. Because firm one's choice is contingent on the state. So if the state is high, we're going to get theta H minus Q2 prime minus Q1 star of theta H.

So you'll see actually here a tricky thing for firm two is these effects are actually going to go in opposite directions. So if firm two-- if it turns out the state is high, that's good for firm one because demand is going to be higher. But it means the opponent is actually going to produce more, and that kind of creates a bit of a those countervailing directions. And then we'll do the other thing. What if what if the state is low? Well, we're going to have to change this and change this. So we have--

Notice a crucial point here. It's the same Q2 prime here as here. Of course, if I were firm two, I would love to choose a different quantity when the state is high than when the state is low. But I'm unable to do that because I'm not able to observe the state. So I choose a single quantity, Q2 prime, and I take expectations over the state.

So here's our conditions. We won't fully solve it out because it gets a bit messy. But let's just reduce each of these to a first-order condition. So here, we've stated it as an argmax property. Let's just write the corresponding equalities.

Well, in this top row, we can use our trick. So I'll just go through these one more time. In the top row, we're going to see that Q1 prime should be half of theta L minus Q2 star of theta L. 2. F1 theta H, we can immediately see that our initial intuitive guess was correct, in that Q1 star of theta H is higher than Q1 star of theta L. Because Q2 star is the same in these two expressions. The only difference is theta L is here, and theta H is here, and we know theta H is higher than theta L.

Now, the harder one, let's say F2. What's the trick here? It turns out we can distribute a little bit. So I guess one nice observation to keep in mind is the half is important to get the right value, but it's not going to affect the maximizer.

Just like up here, if I maximize 3x squared instead of 2x squared, the set of maximizers would have been the same, though the maximum value would have changed. So I don't actually have to really worry about the half.

And then I can factor the Q2 prime, and I think we will get Q2 star equals-- it's going to be an average. So it's going to be maybe theta L plus theta H minus Q1 star theta L I think over four. Make sure I did that right. Yeah, I think that should be good.

And you can immediately make one-- I'm not going to solve-- this is how you got that. I mean, this is algebra. I'm not going to bore you with the details. I guess one nice observation, though, to make sure we see what's going on, is what if we added half of this equation to half of this equation?

If we did that, what we would actually see is it can be shown that, in fact, Q2 star is going to equal Q1 star of theta L plus Q1 star of theta H over 2. So it turns out exactly consistent with our original intuition. The equilibrium quantity that Q2 star is going to choose is exactly going to be half of the quantity that the two firms choose.

And if we did it out, we'd actually find that it turns out the amount Q2 star will actually be-- well, if we did a complete information game where the intercept was the average between theta L and theta H, and we solve for the equilibrium of that game, we'd actually get the same number for Q2 star. But that's just a detail. Any questions on how we solve this? So this is a good thing to know. This will come up on some exams.

So now we're going to change course a little bit and start talking about maybe a more substantive application, really, what I think motivated us to set up Bayesian games, and that's the theory of auctions. So now we're going to start with auctions, and we'll play a little auction game in a second. But I want to start with just some examples of auctions.

Have any of you been to an auction, played an auction, participated or know of auctions? What are some auctions that you've heard about or seen about? Yeah?

STUDENT: Treasury auctions.

IAN BALL: Treasury auctions are a great example. Yeah, maybe one of the biggest ones. So what's a Treasury auction?

STUDENT: When the government basically auctions off government bonds to different retailers, to different institutional investors.

IAN BALL: Great. And this is the way that the price and the interest rates that you see for government debt are set. They're set endogenously by this auction. Yeah. So that's a good example. That's the government. When the government wants to borrow and raise money, they sell debt. And the price that they sell it at is through an auction. Yes?

STUDENT: Fundraiser auctions, where you actually get items.

IAN BALL: Yeah. Yeah, so fundraising auctions. Yeah, what kinds of items are they auctioning off? Yeah.

STUDENT: Cars, vacations.

IAN BALL: Yeah. It's good. Yeah, and I think there's kind of a game element to it, and people show off about how much they bid. And they feel good because they show off that they're rich and charitable at the same time. So that's, I think, why we see auctions used there rather than just prices and saying, make a donation. You to get the warm glow, and you take the car. Any other kind of common auctions? Oh, yeah? Sorry, I didn't see you.

STUDENT: Google Ad space.

IAN BALL: Oh, great. So I think when I've taught this course in the past, no one's given that example, and I always have to bring it up. Google Ads is the biggest-- Google is the biggest private auctioneer in the world. They make $250 billion a year on auctions, roughly, in revenue.

So normally, yeah, every year I ask this question. Everyone always thinks it's like Christie's or Sotheby's or one of these people that's-- one of these firms that's auctioning off paintings. But indeed, Google Ads is a much, much bigger market. So I guess great to see that answer. So how does Google Ads work? Or if you don't, that's OK, but just some thoughts.

STUDENT: Automated at this point. But there's like a second-price auction, where everyone gets on an ad space, and then the top bidder pays the second highest bidder's price. And then it's ranked by the prominence of the ad.

IAN BALL: Yeah, yeah. So you're describing really a lot of details which we're exactly going to go into in a few classes. That's great. So have you-- do you have-- did you work at Google, or--

STUDENT: I'm taking 1419 right now.

IAN BALL: Oh, OK. Great. Good. So a little bit of overlap with that course. So I think Google Ads is two main categories. One is when you search on Google you see those sponsored advertisements at the top. Those are keyword searches where firms are bidding for the right to be displayed. So if you type in vacation, a hotel, or a flight might want to show their ad up there.

The other way they make money is through third-party websites. So if you go to the *New York Times* and you see ads, those ads are also managed through Google. And the *New York Times,* if they want to raise revenue and post ads on their site, they'll use Google as a third party. And then Google will hold an auction on their behalf, where firms might be bidding for the right to show an ad to someone of a certain demographic in a certain place who's doing certain things.

So these ads are very, very targeted. And that's been sort of-- that's where Google makes all their money. You might think, oh, they have a great algorithm for search. But that algorithm just enabled them to get attention, which they then monetized through ads. Yeah?

STUDENT: So when you said about the *New York Times,* so the *New York Times* basically pays Google to take care of their ads through them?

IAN BALL: Oh, Google would pay them. Yeah. So they want to raise revenue on their site. So what would happen is-- I mean, they could try to do it themselves, but that's kind of hard. So they would work together with Google. And then Google is going to conduct these auctions. The advertisers are going to pay Google, and then Google is going to give some cut to the *New York Times* and then keep some cut.

So the advertisers pay. Why do they pay? Because they get to show their ad, and they hope they make money off that. Google is the intermediary that manages this, and they're going to take a cut. And then the *New York Times* is able to monetize some of the attention they generate. Because if the *New York Times* does well, and a lot of people go to their site and look at their ads, then they want to get some money for this. Yeah.

STUDENT: They're like, not *New York Times* ads.

IAN BALL: No, no. What I mean is if you go to the *New York Times,* you'll see ads on the side, which have nothing to do with the *New York Times.* I mean, the *New York Times* might sometimes-- the *New York Times* also shows some of their own ads, but that's just--

And then the *New York Times* would also be a player on other websites. So if you go to the *Washington Post,* you might see a *New York Times* ad next to it. That's because the *New York Times* has bid through Google on an ad to display on the *Washington Post.*

So some firms are playing on both sides of the market. They're both buying ads through Google and also advertising themselves on their site.

STUDENT: So the *New York Times* is getting-- some of the money that's getting paid to Google by the advertisers go to the *New York Times?* [INAUDIBLE]

IAN BALL: Yeah. I mean, I don't know exactly how the accounting-- you could imagine it going different ways. But at the end, the money that's paid by the advertiser-- some is going to go to Google, and some is going to go to the *New York Times.* Yeah. Yeah, exactly. Great.

And then, of course, if it's on Google's own site, then there's no third party, and Google's just taking the money directly if it's a sponsored ad on the site.

I guess the classic one that we often talk about-- paintings, antiques. And I think a basic question we should ask-- I mean, why-- I've given some examples here. Most things we buy in the world are not auctioned off.

When you buy clothes at a store, when you go to a coffee shop, when you watch a movie, when you buy a drink, these things-- generally, there's a price. You get dinner. There's a price that's posted, and you just pay it.

Whereas, in these examples we see auctions used. Do have any thoughts about why we might see auctions in some contexts and prices in another maybe? Yeah?

STUDENT: Price discovery.

IAN BALL: OK. So good word. Yes. Say more about it. Yeah, it's a keyword.

STUDENT: For example, like in a credit auction, like the government's unsure how to value like a certain type of bond. So an auction, there's kind of a bunch of literature that suggests that it's like the most efficient way to both raise revenue for the auctioneer, but then also to efficiently allocate to all the auctions. And it's like a price-revealing mechanism.

IAN BALL: Great. Great. So let's say I have one of a kind antique. It might be pretty hard for me to know how much people are willing to pay for that. So I could just post a price. I say it's $10 million. But I'd worry. What if someone in the audience ended up valuing that a lot? For some reason, they really, really like this antique. Now I've wasted a lot of revenue.

So we see auctions generally when firms are uncertain about how much people are willing to pay. With a coffee, that's a bit different. Dunkin' Donuts sells millions of coffees. They pretty much know what the market price-- they pretty much know what people are willing to pay for coffee.

And also, it would be-- if you had live auctions at Dunkin' Donuts, I think it would be pretty costly. We're talking about pretty small amounts. Though, on the other hand, Google Ads are very small amounts. So maybe there is a way. And maybe eventually, we will see auctions for coffee as everything gets automated, I guess one final point I'll make that I think is kind of a new, exciting area. Google Ads used to be kind of the hot new thing. Google's losing a lot of money on this now because everyone's just using LLMs. And I think there's a lot of interest now in how LLMs are going to monetize this through auctions.

So there's some talk now about Grok or OpenAI might let firms bid for the right to choose the next token that appears. So if you google-- if you ask an LLM, where should I go on vacation? And the LLM's response says oh, a great vacation destination is, you might have firms bidding for the right to select that next token.

And I think we're not quite there. Well, we don't think we're there yet. It's a little unclear, but we may be going that direction. I think that's an interesting domain for future auction work.

So we're going to start with some models of very simple auctions. Obviously, in these applications there's some complexity that we're not going to capture. And we're going to start with a simple sealed bid first-price auction. So we'll play this on MobLab.

What does a sealed bid mean? It means-- in the old days, you had an envelope, and you gave it to the auctioneer that said that was the amount of your bid. So each bidder submits a bid. It's sealed because that bid is not observed.

It's not like what you might see in the movies where people are calling out. That's called an open outcry auction. Oh, I'll bid 100. I'll bid 120, something like that. It's sealed. You're just submitting your bid in an envelope.

First-price-- well, what's going to happen is the highest bidder is going to win, and they're going to pay their bid. So the highest bidder wins the good and pays their bid.

The first-price nature is really referring to the fact that the winner pays their own bid. We'll also talk about auctions where the winner might pay the second highest bid, or something like that. But for now, we're going to just say the winner is going to pay their bid.

What is the private information here? What are these things that-- the information that some players know that other players don't in the context of an auction? Yeah?

STUDENT: How much they value their goods.

IAN BALL: Exactly right. So here, each bidder privately knows what we might call their valuation. The valuation is going to be the most they would be willing to pay for the good.

The way to think about the valuation is if they were to win the auction, when the good, and pay exactly their valuation, they'd be just as well off as if they didn't win the item at all. So it's the amount that they could pay that would make them exactly break even.

So these games, it kind of helps to have MobLab because we have to generate randomly people's valuations. So let's log into MobLab. I'm going to try to set this up. May have been a while, so I'll give you a second. It actually did pretty well. We'll put that on here.

So here, we see the groups. So this is giving us a few different pieces of information here. Let's understand what each of these columns mean. So you were split into six groups. You weren't a multiple of three. So some of you, we added robots to the groups to bid. I'm not exactly sure. They probably follow the equilibrium bidding strategy. I'm not exactly sure how the robots play.

The first column is optimal surplus. What that's referring to is, well, each of the three people in the group have some value for the good, and that's referring to the maximum of those three values. Because the ideal thing, or the efficient thing, would be to give it to the person who values it most, and it's saying how much value would be created.

So here you see these numbers are generally higher than 50 because the maximum of three uniformly distributed numbers tends to be pretty high. And we see the optimal surplus. Then we see the surplus that was actually achieved.

So we see in three of the groups, in groups three through five, we got 100%. What does that mean if we got 100% precisely? What does it mean exactly? Yeah, on the front? Yeah.

STUDENT: Surplus [INAUDIBLE] surplus [INAUDIBLE].

IAN BALL: Yeah. So in words, what's another way of saying what happened in the auction? Who won the auction?

STUDENT: The person who valued it the most.

IAN BALL: The person-- exactly what it means is the person who valued the item the most is the one who won the auction and took the item home. So we see that the allocation was quite efficient.

In one case, we see efficiency was 93%. That means the item did not go to the person who valued it most, but it happened to be that someone else valued it almost as much as the person who valued it most-- 93%. And so not too much of an efficiency loss. In group five, we did a little bit worse-- 82%. And then the last group, I think we just didn't finish there.

And then we can see what the revenue was for the auctioneer. So that's another way of saying, in this case, what is the revenue? Well, the winning bidder pays their bid. So the revenue is exactly the winning bid. And you can see a few differences here.

So let's look at group three. In group three, the revenue is 94, and the surplus was 94, and the optimal surplus was 94. So what happened in group three? Yeah?

STUDENT: The guy with the highest value bid the value.

IAN BALL: Right. And won. And that highest value is 94, right? Exactly. So how did how did that person come-- I don't know. Who was the-- you don't need to out yourself, but it was someone here. And how do you feel about winning at a price of 94? I don't know who this person is. Do you feel like you came away ahead or behind or you won? How are you feeling? Yeah?

STUDENT: Since there was only one period, and I wanted to win the auction, I just bid all I had.

IAN BALL: All you had, meaning your value, or literally--

STUDENT: Like the 94.

IAN BALL: 94. So you bid your valuation. Yeah, I guess-- all you had. I don't know how much have in the game. So this gets at some issues of preferences. If your goal was to win the auction, maybe that was a pretty good strategy. You did win.

But remember, the way that we defined your valuation is that if you win the item and pay your valuation, it's just as good as not winning at all. So in a sense, you won. But in another sense, you didn't do so well because you got something that was worth $94. In the extreme case, think of it as you paid $94 to get $94. So it was kind of a wash for you, but you didn't lose. So that's good.

I think-- let's see how other people did. Who did best? We had the same issue in group 2. The winner bid exactly 59, and I guess didn't do so well. It was only-- in group five they bid 72 when the value was 73. Does that person want to share your thought? OK, yeah?

STUDENT: Well, I thought that that's the max I can bid to at least guarantee that I'll have a positive payoff [INAUDIBLE].

IAN BALL: Yeah.

STUDENT: I mean, there's a chance I still would have won if I did say 70, 71, or something, and I would have a higher payoff. But I think that there's no-- the probability that I win does go lower every time I bid lower. And I still, like even with my bid, even if I do bid, my payoff would be positive [INAUDIBLE].

IAN BALL: OK. And indeed, you did come away with a positive payoff. But I would argue maybe you could have gotten an even more positive payoff if you bid a little bit lower. Let's see. What about the person-- someone bid 40 and won and valued at 44. So I think they were the overall winner. They had the most surplus of anyone. Yeah?

STUDENT: [INAUDIBLE] my valuation is pretty low. So I thought I would lose.

IAN BALL: OK.

STUDENT: That seemed like a reasonable amount that [? I've got ?] to lose.

IAN BALL: So some people would argue, OK, I think I'm going to lose. My value is low. Why don't I bid more than my value to have some chance of winning? A lot of people reason that way. You didn't do that. I think that would have been a bad idea. Why did you not? You thought you were going to lose. Why didn't you bid a lot more to try to win?

STUDENT: Because then my-- if I did win, my valuation would be negative. So I thought, in this case, I would rather get a valuation of 4 or 0.

IAN BALL: Exactly. Right. So even if your valuation is really low, it's true that you're unlikely to win. But that might tempt you to overbid. That's not a good idea. Because if you do win, you're actually a loser because your valuation is smaller than the amount that you bid. So you bid $4 less than your true valuation. In the end, you got-- well, I don't know who else was in your group, but you got-- there was someone who had a value of 47. So actually you got lucky, and you beat someone who actually had a valuation higher than you, and you came away with a surplus.

So my overall, I would say, people are maybe not ready to bid for Google or bid for the *New York Times.* I think there was some poor bidding-- I don't know-- OK bidding here, I think people generally bid a bit too high.

So if you see they've actually computed the Nash equilibrium line, which is the dotted green line. And you can see everyone's actually bid-- I think every single person bid higher than would be optimal to bid. And I think this is a very common finding when you do things in the lab.

And I think it gets to, do we have the right model of preferences? I think there's a certain excitement people get for winning. And they want to win, even if it means they end up paying a little bit too much.

So people tend to overbid in first-price auctions. And that might actually be a reason why we see a lot of first-price auctions in practice, even though the theory would tell us that second-price auctions have some more desirable properties that we'll talk about.

Maybe let's do it one more time, but now let's make it a second-price auction and see what happens. How do I change it? I'm not going to have it. Now, I just have to add it again. That's silly. So now let's make it a second-price auction.

So what does a second-price auction mean? Well, the highest bidder is still going to win the item, but now they're going to pay the second highest bid rather than their own bid. So this should change how you play. So curious how things went.

So here, let's look at group one. Here, we actually were in groups of six. So the optimal surpluses tended to be very high because we're looking at the maximum of six different random variables. It was efficiently allocated every single time. And it looks like the revenue was-- I guess it's a bit hard to know from the data we're seeing here, but it looks pretty good.

It actually looks like people are bidding basically their valuations. This is basically a straight line. So people were actually very, very consistent. Do people want to share how they thought about this? I'm actually a bit surprised by this. This is not how it usually goes. Yeah, in the back?

STUDENT: I think in my notes there was like an analysis of second-price auctions. And it says like it's most optimal to always bid [INAUDIBLE].

IAN BALL: OK. That's a good answer. Yeah, that's exactly right. So maybe people read the notes. I'm surprised so many people have read the notes. I thought-- so I think I've done this in past years.

And normally when I tell people they bid too high in the first-price auction, we then go to this, and they bid really low. But indeed, if it's a second-price auction, it's actually optimal to bid your true valuation and not to bid less than your true valuation.

What's different here? I said before, if you bid your valuation, then you can never win any money. So what's different in the second-price auction? Yeah?

STUDENT: [INAUDIBLE] my valuation. I think someone said earlier, [INAUDIBLE] highest. I've got a pretty high chance to win, but then I'm actually getting a price that's below that because you're [? being ?] the second highest bidder. So you know you're getting a positive value back as well.

IAN BALL: Exactly right. So in the first-price auction if I bid my valuation, then if I win the auction I'm actually going to pay my bid. So I'm going to pay my valuation. And therefore, I don't come out ahead.

But in the second-price auction, if I bid my valuation, well, I'm not actually going to have to pay my valuation if I win. I don't know what I'll pay, but I'll pay the second highest bid, which will always be strictly lower than my valuation. And we'll go over the argument, I think, next class about exactly why bidding your own value is exactly optimal in a second-price auction.

And this is why one nice thing about second-price auctions is you don't actually have to think too much about how much other people value the good. In the first-price auction, you really have to think hard about how much the other people are going to bid, and try to respond strategically to that. So let's close this down. And let's actually analyze the simplest possible example of a first-price sealed bid auction and try to solve for the equilibrium of this auction.

So solving for auction equilibrium can get a bit messy. So let's start with the simplest possible example, which is a first-price sealed bid auction with two bidders. i equals 1 and 2. And let's assume that each bidder's valuation for the good is uniformly drawn from between 0 and 1.

So before in the example, in the MobLab example, I think it 0 to 100. We're just going to change units now. So bidder i's valuation, we'll call that Vi-- So bidder i is valuation Vi. This is going to be in 01, and it's uniformly distributed and independent.

Uniform just means that the probability that my value is in any subinterval of 0, 1 is just the length of that interval. So the probability that my value is between 0 and 0.2 is just 0.2. So it's pretty easy to calculate these probabilities with uniform distributions. Another way of saying that if you've taken some probability is the PDF, the probability density function is just 1.

Independent-- what I mean by that is that bitter one's valuation is independent statistically of bidder two's valuation. So what that means is if I know my own valuation, V1, I still believe that my opponent's valuation is uniformly distributed. My valuation provides no information about the valuation of the other bidders. At the extreme, other case would be we all valued exactly the same amount. But that's not what we're studying here

And each bidder can submit a bid. We'll call it bid. Let's say bi, which is just a non-negative number simultaneously. So here are the bidders. They each privately draw their valuations. Each bidder knows her own valuation, but not the other player's valuation, and then submits the bid, bi, to the auctioneer.

And then what happens? Well, the higher bid, the higher bidder, that is, the bidder who bids more, wins good, wins the item, and pays bid.

So I think the first step to analyzing this auction is to try to understand what each player's payoff is in the auction. What is their utility as a function of their valuation and how much they bid? So let's go over here.

So formally, this action induces a Bayesian game, where each bidder is-- the action they can choose is how much to bid, and their payoffs are going to be what we write down here. So let's look at player one. We could reason about player i, but I think that just makes it a little more abstract.

So let's look at U1 player one's utility as a function of well, what are the things that can affect player one's utility? How much they bid, how much the other player bids, and also how much they value the good. So we're going to write U1 of b1, b2, v1. So these are the bids of both players.

Why does the other player's bid affect my payoff? So I'm player one. Why do I care about the other player's bid? Yeah?

STUDENT: Because if you you've been higher than the other player, then you have to actually consider your value in finding the payoff. Otherwise, that other player bids higher than your value [INAUDIBLE], or your utilities.

IAN BALL: Right. So the other bidder's bid affects me because it affects whether I win. It doesn't affect how much I pay if I win. I always pay my own bid. But it affects whether I win. So that's relevant to my utility.

Also, of course, my valuation, v1, is also relevant. This is bidder one's valuation. It turns out bidder two's valuation doesn't enter here because all I care about is what we bid and how much I value the good. Of course, in the game, the other player's valuation might affect how much they choose to bid under their equilibrium strategy. But in terms of just the payoffs, we just want to write payoffs as a function of actions and types, or state. So let's try to fill this in.

So if we want to be really formal, then v1 would be playing the role of your type or of the state of nature, going back to our formalism about Bayesian games. So now I want to write down this payoff function. I think it's helpful to separate into three cases.

And in a sense, I should say that what we're going to fill in here is basically what defines the first-price auction. The rules of the first-price auction are going to determine what the payoffs are here. If we had a different auction format, then our payoffs, as a function of the bids, would be different because those bids would be aggregated in a different way.

So let's say b1 is less than b2. I think this is the easiest case. What is player one's utility going to be in this case. Zero, right? If their bid's lower they lose. If they lose, they win nothing. They pay nothing. The utility is zero.

Let's now look at this case. What's their utility going to be here? Yeah?

STUDENT: b1 minus b1.

IAN BALL: Exactly. So what happens here-- I'm bidder one. I bid higher than my opponent. That means I'm the highest bidder, so I win. Because I win the good, I get my valuation v1, but I have to pay my bid. And my bid was b1. So v and b sound kind of similar, but my total payoff is v1 minus b1.

What about this case here? Well, we have to make a choice about how we break ties. Normally, when things are continuous, this isn't so important. Normally, what we imagine is if both people bid the same amount, we flip a coin to see who wins.

So we flip a coin. If it's heads, player one wins and pays their bid. If it's tails, player two wins and pays their bid. So what is player one's expected payoff going to be in this case?

STUDENT: v1 minus b1.

IAN BALL: Exactly right. So there's a half chance they win the coin flip. If they lose the coin flip, they get 0. If they win the coin flip, they get this payoff. So we just get v1 minus b1.

STUDENT: So like, how we model it isn't-- probably doesn't matter too much [INAUDIBLE].

IAN BALL: Exactly right. Exactly right. Let's just, as an exercise, just to make sure we understand things we're going to solve the first-price auction. But let's just imagine if it was a second-price auction, let's see how these payoffs would change, so just to make sure we're on the same page.

If it were second-price, well, the number of bidders are the same. The valuations are the same. They're still bidding as before, except instead, the higher bidder wins the item and pays the second highest bid. Then we're going to have a different expression for this utility. What would it be? So let's just go through on the other side. I think this is still the easy case. What if this is second-price auction, then what's the payoff here for bidder one?

STUDENT: We would have b2 instead of b1.

IAN BALL: Exactly. Yeah. So otherwise, everything is exactly the same. Good. So it's still 0 here because if I bid less than the other player, I lose. Let's go down here. If I bid more than the other player, I win. But now I play the second highest bid, not the highest bid. If I'm bidder one and I bid more than the other player, then their bid is the second highest bid. So I'm going to get v1 minus b2.

And indeed, in the middle we're going to flip a coin again. So there's a half chance I win, in which case I get v1 minus-- well, here it actually-- notice, in this case b1 equals b2. So I could write b1 or b2. It won't matter, but I'll write b2 because I think it's more intuitive. So exactly right.

So you'll see how this is going to fundamentally change the incentives of the bidder. A key difference, a special property of the second-price auction, notice, is that your bid b1 affects whether you win, but not your payoff if you do win. In a first-price auction My bid b1 Affects whether I win, but also what my utility is if I do win. And these two counteracting forces. So now let's see if we can solve for an equilibrium.

So let's go through our steps. The first step is to figure out what kind of objects the strategies are. We first have to figure out what the strategies are. So let's say solving for Bayes-Nash equilibrium. What is a strategy for-- let's say we have firm one and firm two.

So as before, a strategy has to specify what action you're going to take as a function of your private information, as a function of what you know because you can only condition your action on things you know. So what can firm one condition their bid on? Their own valuation.

And indeed, that's what a strategy is going to be. It's going to have to specify my action. That is, how much I bid as a function of my valuation. And this is where understanding function notation is really important. We have to careful about this. So it's going to be a function b1 from 0, 1 to-- maybe I'll write R plus to denote the non-negative real numbers.

So what does this say? My strategy is a function, b1, that takes in any possible valuation I could have for the good. We said the valuations live in 01, and it assigns to each of those valuations an amount that I bid.

So what's tricky here is that before, in our Cournot example, firm one only had two types. So we had to specify two numbers-- how much they produced in the high state, and how much they produced in the low state.

Now we have a continuum. So one way of thinking about a function is it's like a continuous vector. We're having to specify infinitely many, in fact, continuously many numbers. And that's why solving for auction equilibria is going to be quite hard. Because we need to understand not just how each bit are actually bids, but how they would have bid if their valuation were different. And we're going to have the same thing for two.

So we've got to solve for a lot of things here. We could write out the full equilibrium condition. But I think the trick with a lot of these auction equilibria, at least in this case, is we're going to make a guess. We're going to use guess and check.

So what's our strategy? I shouldn't say strategy. That would cause confusion. We want the guess and check. What's our approach? So we're going to look for an equilibrium of a very specific form. So we're going to guess a symmetric equilibrium. I think it's pretty intuitive that the equilibrium will be symmetric here because the distributions of the valuations are symmetric.

Now remember, symmetric doesn't mean the two players will always bid the same amount because the two players may have different valuations for the good. What symmetric means is the two players use the same strategy, which means if player two had a valuation of 0.7, she would bid the same amount as player one would bid if she had a valuation of 0.7.

So in a symmetric equilibrium, b1 of v-- I'll just use v to denote a generic value-- equals b2 of v for all v. So this says if bidder one's valuation is v, the amount that she bids is the same that bidder two bids if her valuation is v. Maybe I'll put stars here to show that this is an equilibrium.

But we want to guess a bit more about the structure of these bidding functions because now we have to solve for infinitely many numbers. That's a bit hard. So we want to guess that the bidding function takes kind of a nice form. And it turns out that the trick here is to guess what's called a linear bidding function, where everyone just bids their value times a fraction, times a constant. So let's call it alpha v.

So we're guessing a symmetric equilibrium that takes this special form. And we're going to try to solve for alpha. Any guesses about the value that alpha might have? Is it going to be bigger than 1? Less than 1? Positive, negative? Any thoughts about this thing, alpha? Yeah?

STUDENT: Something positive less than 1.

IAN BALL: Positive less than 1 is a good guess. Because we know, and we saw from the example, that I want to what's called shade my bid. I want to bid less than my true valuation. So certainly, alpha should be less than 1.

You can't bid negative amounts. And we'd like your bid to be-- we think it should be increasing in your value. So it makes sense for alpha to be non-negative, or it's really going to be strictly positive. So in general, the idea is you're going to get some value v. And instead of bidding your true value, you're going to bid some fraction of that value. And our goal is to solve for that alpha. So step one is our guess. And step two is we're going to check that this is actually an equilibrium.

Now the checking is the hardest-- is a tricky part. And I want to make-- there's a really important point here. It's not enough to say suppose my opponent is using a linear bidding strategy, alpha v, then the best linear bidding strategy for me is also alpha v. That's not enough because I could deviate to a nonlinear strategy.

So to understand, we're looking for an equilibrium where both firms use linear strategies. But the firms could deviate to nonlinear strategies as well. We're not restricting them to using linear strategies. So the way we check this is we want to say-- so check. if b2 of v equals alpha of v. So if bidder two is following this strategy, then the best bidding strategy for firm one, or for bidder one, is also alpha v.

So we're going to suppose that my opponent is using this strategy and then compute, what's the best way for bidder one to bid and check if that happens to be the same strategy. If it doesn't work out, then that means our original guess was wrong, and there is an equilibrium of this form. If it does work out, then we'll be able to solve for alpha.

So I guess that's a key step of guess and check. The output of guess and check will either be I made a good guess. Here's the equilibrium. Or it will be my guess was wrong. I guessed an equilibrium of a particular form that didn't exist, and then let me try again. Of course, in class, we're going to make good guesses. But on your homework, you might make a guess that it turns out to be wrong. OK, let's move over here.

So following our guess and check strategy, we're going to suppose b2 of v equals alpha v for some alpha that we're going to try to solve for later, so for an arbitrary alpha. Let's say for some fixed alpha that's maybe between 0 and 1.

And now we want to consider bidder one with valuation v1. And we want to say, put yourselves in the shoes of bidder 1. Suppose your valuation for the item is v1. Suppose you believe that the other bidder is following this bidding strategy. What is the optimal bid for you as bidder one?

So what are they trying to do? Their goal is to maximize over b1 their expected payoff. Well, let's go over and look at our formula and try to compute our expected payoff.

Now, the tricky thing is we can't just plug in a single number for firm two's bid, bidder two's bid. Because bidder two's bid depends on bidder two's valuation, and we don't know what that is. So what we're going to have to do is take expectations over what the value of bidder two is. So we're going to maximize what? We're going to take expectations of u1, b1, b2 of v2, v1. I think this is the key step. So let's understand what's going on here.

What I'm taking expectations over is v2. This is the random thing. Because from my perspective as player one, I don't know what the valuation of bidder two is. So this is the thing I don't know about.

v1 is not random. v1 is my valuation is bidder one, and I know what that is. So this is fixed. b2 of v2 is the amount that bidder two bids when their valuation is v2. And under our supposition, we know that this is going to be alpha v2.

But even though we know what bidder two bids as a function of their valuation, we don't know what they actually bid because we don't know what the valuation is. So it's alpha v2, but we don't know what v2 is. And then b1 is a number. This is my bid.

So just to be clear, what's inside the expectation is, what is my payoff going to be if it turns out that the valuation of bidder 2 is v2? And that's what we get in here. And then I'm taking expectations over that. So now let's try to simplify this.

Well, in a lot of cases, this is just zero. This is only going to be non-zero when my bid is bigger than my opponent's bid. And when my bid is bigger than my opponent's bid, this is exactly what I get. So it turns out we can write this as this. Let's understand what's going on here.

To compute my expected payoff, we just have to say, well, there's two possibilities. Either I win, or I lose. If I lose the auction, I pay nothing, and I get nothing. So my payoff is zero. So we don't even have to worry about that. You can think of it as adding plus 0 times the probability that my bid is lower than the other if you want, but that doesn't change things.

v1 minus b1 is what I get if I win, but I have to multiply that by the probability that I do win. One second. What is the probability that I do win? Well, I'm going to win if my bid, b1, is higher than my opponent's bid, alpha v2.

And we really have to understand crucially here in this probability, the random thing is v2. b1 is not random. b1 is a number. I know what b1 is. What's random is v2, and that's what we're taking probabilities over. Yes?

STUDENT: Are we just not considering the type case because we have a continuous distribution to the probability or that?

IAN BALL: Exactly right. So you could add a third-- really, we could think of it as three numbers. It's this times this probability-- maybe we'll write it out to make it clearer-- plus 0 times the probability that b1 is less than alpha v2, plus v1 minus b1 plus 1/2 the probability that b1 equals alpha v2. This is the full expression.

But this disappears because we have a 0. And this disappears because the probability that b1 is exactly alpha v2 is 0. Why? It's always tricky to say the probability that b1 is something. Remember, b1 is just a number. So what I really mean is the probability that alpha v2 is equal to this fixed number b1. That's also 0. So that cancels, and I just get this expression.

Now we see the basic trade-off that we always face in auction theory. If I increase b1, then when I win the item I pay more. And therefore, I win less when I win, but I win more often. If I reduce my bid, then winning is even better because I pay less when I win. But winning happens less often because I'm less likely to win. And the optimal bid is exactly going to trade off those two forces.

Notice, I'm not putting weight on just whether I win or not. It's how well I do when I win. If my goal was just to win, or just to win and make positive profits, then I would just bid exactly or just below my valuation. But I care about how much I win when I do win. So let's go to a new board and finish this off.

Well, let's write this in a different way. We have v1 minus b1 times the probability-- well, the random thing is v2. So let's focus on v2. This is v2 less than b1 over alpha. But what is this going to equal for a uniform distribution? Yeah?

STUDENT: b1 over alpha.

IAN BALL: Just b1 over alpha. We have to be a little careful that if b1 over alpha is-- if b1 is bigger than alpha, then it could be 1. But it turns out we usually aren't that careful about that. So let's just say if b1 less than or equal to alpha.

It turns out you're never going to want to bid more than alpha. Because if you bid more than alpha, you're already going to win anyway, so you're just increasing the amount that you pay. Let's understand.

If my opponent's strategy is alpha times v2, the highest they could possibly bid is alpha. So it would be crazy for me to bid more than alpha because I'm already guaranteed to win by bidding alpha. So it's a safe assumption that b1 is less than alpha. And in that case, then we indeed get b1 over alpha.

So now if we simplify, we get v1 minus b1 times b1 over alpha. And if we maximize this, we're actually going to get b1 star equals v1 over 2. Because the alpha doesn't affect-- you could do the first-order condition. But it's our classic trick-- b1 times v1 minus b1 is maximized by b1 equals v1 over 2.

Wait a second. Was our guess wrong? We guessed that player two was using a bidding strategy of alpha v. And our best response for bidder two is to always bid v1 over 2. So was our guess right? How can we make sure that this agrees with our guess? How do we have a symmetric equilibrium? Alpha equals 1/2, right?

So we conclude that alpha equals f. And let's just go through the reasoning right. If alpha equals 1/2, then both bidders are using the bidding strategy where they bid exactly half their value. Given that the other bidder is bidding half their value, well that's a special case of alpha times their value. So my best response is half my value. And that shows that I'm bidding optimally given the way my opponent is bidding.

So the key property is that when we solve for the maximizer, we got something within the linear family. We got a new linear function with coefficient 1/2, and then we just matched our coefficients. And we said alpha equals 1/2. So we'll continue this on Thursday. Great.