[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: And today, we're going to discuss a new topic, which is bargaining with incomplete information. Incomplete information. And you recall that earlier in the course, we studied a model of bargaining, but in that model, we had complete information. And when we studied bargaining with complete information, we observed-- so let's recall that with complete info, we had no disagreement and no delay.

So when the two sides were bargaining, whenever there was an agreement to be reached, whenever there was some possibility for agreement, it was always achieved, and it was always achieved in the first period. There was never any of this costly delay.

But of course, in reality, we see a lot of disagreement and a lot of delay. When we look at real bargaining in practice. And we argued that maybe the missing piece of our previous model was this very strong assumption of complete information, the assumption that both sides exactly understand the preferences of the other side.

And just to make this a little more concrete, let's look at an example-- let's say plea bargaining. So this is common. You might have a defendant who's accused of a crime and a prosecutor. And before potentially going to trial, they might negotiate some plea agreement which says you don't go to trial, and instead, you serve a certain amount of time.

And here, I think there's definitely a role for incomplete information. So let's put yourself in the role of the defendant, and let's say you really want to reach a plea deal, you don't want to go to court, so you make a really generous offer to the prosecutor. You say, look, I'll plead guilty and I'll serve 30 years to avoid going to court.

Now generally you would think making a very generous offer would help you get the offer approved, but what might the prosecutor think if you say, I'm happy to serve 30 years? Yes? That you're super guilty or maybe the evidence is really-- yeah, the evidence is really, really strong, it's going to be easy to show in court.

So there's this perverse effect where normally a better offer is more likely to be accepted, but now a better offer might suggest, wait a second, there's a reason she's giving me a really good offer, and maybe I actually don't want to accept it and I'd rather go to court.

So in this case, the defendant has private information about certainly their guilt, but maybe more relevantly, how easy it is to prove their guilt. And the defendant, their offer maybe conveys information about-- we'll just say guilt or evidence of guilt or other things.

Now let's look at the other side. The prosecutor may also have private information. Let's say the prosecutor also really wants to get a deal. And the prosecutor says, look, I'll give you a really generous deal, you only have to serve six months in prison if you plead guilty, and then we don't have to go to court. That sounds like a great deal. Why might the defendant be hesitant if they received an offer like that? A really generous offer by the prosecutor? Yeah?

AUDIENCE: If we assume the prosecutor doesn't have any evidence, so if they go to court, the defendant might be able to have--

IAN BALL: Exactly. So this is, I think, the key question. They don't just say how good is the offer, they say, what is the quality of the offer? Tell me about the other side's information. And maybe they just have no case. So maybe I should go to court even though this offer seems really good.

So here, the prosecutor's offer might convey info about maybe the evidence they have. And I think-- legal scholars certainly think this is exactly why we often see disagreement and delay in bargaining because the kinds of offers that you might make to try to reach agreement can convey information that make it harder to reach an agreement.

So today, we're going to look at two general classes of bargaining games. And I think it's helpful, at first, to distinguish two different situations, signaling and screening. So, so far, we already talked about signaling games last class where someone might choose a level of education to try to signal their ability as a worker. And more generally, in signaling games, what distinguishes them is that the first mover has private information.

So when the first mover, the player who takes an action first, has private information, their action is going to convey something about their private information. So their action might reveal or convey information.

And when the first mover is strategic, they're going to think through how their action might convey information, and that's going to affect their choice of action. In screening. It's the second mover that has private information. So you might say, oh, well the first mover, they're going first, they don't have any private information. Private information doesn't have much of a role.

But it still has a role because what the first mover is potentially going to do is they're going to choose their action, anticipating that the response to their action conveys information that the other side has. So if I make an opening offer to you, and your choice whether to accept or reject that opening offer conveys to me information that you have.

So here, they're signaling information here, you might call they're kind of fishing for information. They're choosing an action so that the response to that action conveys information from the other party. So maybe I'll say fishing for information. And the formal term screening-- oh, I guess the first mover, I should say, is the one.

And the term screening comes with the idea that your response-- I'm the first mover, your response to my action screens you and distinguishes different players with different private information. So it's a screening mechanism. I might offer something to you. I might say, oh, you really want this? Well, let me offer it to you for a high price, and that will screen out people who actually want it and are willing to pay a high price from people who are not. That's where the term screening comes from.

So today we're going to study bargaining with incomplete information, we're going to look at both the cases of signaling and screening, and we're going to do it all within a canonical framework of just a buyer and seller negotiating over the price of a good. Of course, these ideas can be applied much more generally, but these games are a bit harder to analyze than the ones we've done before, so we want to stick to a very simple model to begin with, but the lessons extend a lot more broadly.

So today, we're going to look at a buyer and a seller. And they're going to be bargaining over the price at which they exchange the good. So we have our buyer and our seller. The buyer has some value v for the good. And we're going to make different assumptions about who knows v and who doesn't, but it's always going to be called v. So this is how much the buyer is willing to pay for the good.

And the seller has some cost c. I should point out here that it's just nice to use v and c And think of c as a cost, but another interpretation of the cost of the seller is how much they value the item. Because remember, the cost c is what the seller loses when they give up the item. And we can interpret that utility loss in two ways. Either they actually have to produce the item at a cost in order to give it to the other player, or they actually like it themselves, and when they give it up, they lose that.

So we'll call it a cost c, but you could also think of it as the seller's value, it's just, this way, when I say value, we know we're talking about the buyer.

So what are their utilities? Well, if they exchange at a price of p, then the buyer's utility is v minus p. They get value v from the good because they're exchanging, but they have to pay a price p to the seller. And then the seller's utility is p minus c because the price p is what goes to the seller, the seller gets p dollars, but has to bear the cost c either of producing the good or of losing it.

And then if there's no exchange, well now, we just say each player gets utility 0 because the good is not exchanged and money does not change hands. So here, we have uB equals 0 and uS equals 0.

Now let's-- to make some progress here, at first, let's first think about, well, if we think we're going to-- if we might see exchange, what are the possible ranges of prices at which exchange takes place? Well, at least if v and c are known-- so let's start with this case. So case 1 is kind of the easy case. Case 1 is where v and c are publicly known. So this is the case that we studied earlier in the course. This is actually a case of complete information.

So let's review this, and then we'll contrast it with the other cases. Well, if v and c are public, there's something we can say. If trade happens at p, we can say a lot about what p must satisfy. If v and c are public-- let's suppose we see the players trade at price p. What do we know about how p compares to v and c? Yeah?

AUDIENCE: So p has to be greater than c, but less than v.

IAN BALL: OK. And can you explain why? Yeah?

AUDIENCE: The reason is that we need the p to be at least 0, otherwise the buyer would choose not to exchange.

IAN BALL: Exactly right.

AUDIENCE: --p minus c could be greater than or equal to 0 for the seller to want to exchange.

IAN BALL: Exactly right. So we have two constraints. If they exchange at price p, this is voluntary exchange. It only happens if both players agree to it because either player can walk away and say no. So the only possible prices we could exchange are prices p that make each player weakly better off than their outside option of just walking away.

So if we trade at price p, we must immediately have-- maybe I'll write it in two separate inequalities-- p greater than or equal to c and p less than or equal to v. This inequality is because of the seller, so maybe I'll put S here. If p were less than c, the seller would never want to trade at price p. And this comes from the buyer. If the price were greater than v, the buyer would simply walk away.

So we immediately see from this that if trade is going to occur, we must have c smaller than p smaller than v, and in particular, that means c must be smaller than v. So let's assume that c is strictly smaller than v. This is a common assumption to rule out the trivial case. If v is strictly smaller than c, then there's no way we can ever have exchange. And if v equals c, well, then we could have exchange, but it's only going to be at the price p equals c equals v, and then there's no change in utility and it's not a very interesting case.

So this is where sometimes what we say here is there are gains from trade. Some people say this is what economics is all about, gains from trade. Why are the gains from trade here? Well, when the good moves from the seller to the buyer, the buyer's utility goes up by v and the seller's utility goes down by c. And because v is bigger than c, it's possible for there to be a price such that both players are better off for. The gain to the buyer is bigger than the loss to the seller when we exchange.

So in this case, we've done it before, but let's just go through the analysis. We need to make a choice, our modeling choice about the order in which the players move. So I think it's pretty standard, at least in practice, when we see-- certainly in stores, we see Sellers often post a price first. You could maybe walk in and offer your own price, and in one-off negotiations, that does happen, but let's think of the case where we have many goods, we tend to see it done this way.

So let's make our timing be that the seller proposes a price. And then the buyer-- of course, it's always the potential buyer, they don't have to actually buy-- observes the price p that's offered, and then chooses whether to accept or not. Accept the offer or not, meaning buy or not buy. So maybe we'll say buy or not buy.

And if we want to draw this in extensive form just as a-- so maybe this is a good review as we're getting closer to the exam. So we have the seller here. Now it's a little tricky because the price can take any value, say any non-negative value. So there's infinitely many prices the seller can offer. It's always a little hard to write that down, but maybe we'll write it something like this. I'll call this the price p. I'm visualizing this as the full range of prices the seller can offer. Maybe I'll put a distinguished one here.

And then the buyer sees the price and makes a binary choice, I'll call it A and R, accept or reject. Accept means buy, reject means not buy. And then let's make sure we get the payoffs here. Remember, v and c are known, these are just parameters of the game. So if the price p is accepted by the buyer, then I think we put buyer first-- so let's do buyer utilities first. The buyer is going to get v minus p, the seller is going to get p minus c. And over here, it's just going to be 0, 0 because both sides reject.

So now, let's try to understand, let's try to solve for a subgame perfect Nash equilibrium of this game. Notice, we can use the solution concept of subgame perfect Nash because there's no private information because this is a game of complete information. Once we move to incomplete information, we have to use perfect Bayesian equilibrium, but for this, we can just use subgame perfect Nash. So let's look for subgame perfect Nash equilibrium. What we want to do?

Well, the trick with a game like this is we always work backwards. In fact with this game, we can apply backward induction. So let's start with the second player. This is the buyer. And we can ask, well, what offers is the buyer going to accept? Well, the buyer can either get 0 or take the offer. So they're going to accept an offer if their total utility from accepting the offer is at least 0. So the buyer is going to accept-- maybe I'll use a double arrow to say if and only if v minus p is greater than or equal to 0.

This is the utility of accepting the offer, this is their utility of rejecting the offer, and they're going to accept if the utility from accepting is better than the utility of rejecting. There's a little flexibility if this equals exactly 0, but a common theme, when you're looking for equilibria, it's often better to break ties this way and say that when indifferent, they accept. It tends to make sure the sender deals with some existence problems for the sender's choice.

Well, another way of saying this is they're going to accept if and only if v is greater than or equal to p. And anticipating that, we can now go to the first mover, the sender-- the seller. And what is the seller going to do? If they anticipate the buyer will accept if and only if v is greater than or equal to p, well, then, what price do they want to offer?

AUDIENCE: Well, wouldn't you just offer the highest price possible?

IAN BALL: The highest price possible that will be accepted. And in this case, that's just p equals v. Notice here, we're using our assumption that c is less than v. Because if v was less than c, they might not actually want to offer v because that wouldn't cover their costs, but under our assumption that v is at least as big as c, then the seller is going to make-- offer the highest price that will be accepted, that's p equals v, and then we can see what are the utilities going to be.

Well, the buyer is going to get v minus v. The first v is their value. The second v is actually the price, but it equals the value, so that's going to be 0. And the sender is going to get v minus c. v here is playing the role of the price minus the cost, and this is strictly greater than 0.

Sometimes terminology you might say, remember, we called v-- v minus c, we said there's gains from trade because v is greater than c. We sometimes think of v minus c as the potential gains from trade. This is the amount of utility that can go up if they trade. And here, we sometimes say the seller has fully extracted the gains from trade. Maybe I'll say this-- or maybe I'll say the seller gets all gains from trade.

In this game, we can think of this-- remember, before we studied this game, split the dollar. v minus c is basically the dollar. This is the potential gains that are on the table. It's going to be split in some way between the buyer and the seller, and in this case, the buyer gets none of it and the seller gets all of this. And this is because the seller gets to move first, anticipating the way the buyer is going to respond.

So now that we've reviewed this case, let's now introduce incomplete information. So let's make different assumptions about v and c, unless there's any questions on this. So we're going to start with-- I mentioned over here, there's two cases, signaling and screening. And let's be really clear here. This game didn't involve any signaling or screening because the seller and the buyer-- or in this case, the seller didn't have any private information. So the price they offered didn't convey anything, and this didn't fit into our framework of signaling and screening, but now we're going to move into that case.

So let's start with screening. So what we want to imagine is that the buyer-- we're still going to assume that the seller makes an offer, but now we're going to assume that the buyer has private information. Buyer, but not seller.

So let's be a little more precise. The seller doesn't have any private information, so let's assume that the cost c of the seller is publicly known. This is commonly known, everyone observes this. So we'll say the cost c is known-- or is public. What's not known is the buyer's valuation, how much the buyer values the good.

But v is private information of the buyer. So this fits into our screening setting because it's the second mover, the buyer who has the private information. The first mover, the seller, doesn't have any private information. But as we said before, the second mover's response to the first mover's action is going to convey information.

In this case, the first mover's action by the seller is to offer a price. The buyer, then, is going to choose whether or not to buy at that price. And that's going to reveal information about the buyer, and in particular, it's going to screen or separate different kinds of buyers because buyers who value the good more are going to choose to accept the offer and buyers who value the good less are going to choose not to accept it.

So let's try to model this. Well, If c is known, that's just a number that everyone knows. We don't have to say anything more about that. But v, because it's private information of the buyer, we have to specify the seller's beliefs about this quantity v. So let's make our standard assumption. Let's say that v is uniformly distributed-- this just makes the algebra a bit easier-- is uniform over 0, 1. It's a random variable, it's uniformly distributed over 0, 1.

And now to make things interesting, let's assume that c, this public number, is somewhere between 0 and 1. Because if c was greater than 1, well, we'd never have trade, and if c was less than 0, well, we'd always have trade, it's not so interesting. So let's assume c greater than 0 is known.

And I want to highlight this. I think it's often a source of confusion, when people see a variable, a letter, they think somehow it's not known. It's just playing the role of a number. Everyone knows it. Maybe it's 0.3, maybe it's 0.7, but we're just using a abstract symbol c to make it things a little bit cleaner. Yes?

AUDIENCE: Does the seller note the distribution of v?

IAN BALL: Yes. So when we say that v is uniform, everything we're saying here is common knowledge among the players. So not only is it distributed according to the uniform distribution, but if we're really precise, the seller knows that, the buyer knows that the seller knows that, and so on, everything I've said here is common knowledge. Yeah, good question OK. Great.

So now let's-- ah, c greater than 0, and I also see-- let me be more precise, it's 0 less than c less than 1. OK, so now let's look at the seller's problem. Let's suppose the seller offers a price of p. Well, if they offer a price of p, who's going to buy? Well, the buyer is going to buy if and only if their value is at least p. So then buyer accepts or buys-- kind of use those interchangeably-- if and only if v is greater than or equal to p.

And let's be a little more precise here. What kind of price might the seller reasonably offer? It doesn't make sense for them to offer more than 1. If the price they charge is more than 1, it's never going to be accepted. And if it's less than 0, well, that means they're paying you to get the good, that doesn't make sense.

So let's focus on a price offer that's between 0 and 1. And we can show that they would never want to offer a price outside that range. Now, the buyer is going to accept or buy-- slash, buy-- if and only if v is greater than or equal to p. Well, let's think about how likely this is, because as the seller, all I care about is when I offer you a price of p, what is the probability that you actually buy at that price?

So we can think-- we can draw a little graph-- here's 0, here's 1, the buyer's value is uniformly distributed over this interval, the seller is choosing some price p in this interval, and the buyer is going to buy if and only if their valuation lies above p. So lies in this segment.

Because we're working with the uniform distribution, the probability that the buyer's value lies in this segment is just the length of the segment. That's a nice, convenient property of the uniform. So the probability here is just 1 minus p.

So now, we can compute what is the seller's profit from offering a price of p? Well, as usual, profit is the number or the probability that you sell the item times the profit you get when you do sell it. So if they do sell it, their profit is going to be p minus c because the buyer buys at price p, pays p to the seller, but remember, the seller has to pay a cost c to sell it.

But we have to multiply that by the probability that the buyer actually sells it-- or that the buyer actually buys it, sorry, which is, here, the probability that v is greater than or equal to p. And from what we said down here, this is just 1 minus p.

So as usual, we have the two counteracting forces. When we raise the price p, "we" being the seller, we earn more per item, but we reduce the probability that the item is purchased. And that's the fundamental trade-off. If you set the price really high, it's great when someone buys, but they're not very likely to buy. And if we just use calculus or use our trick here, we can see that the solution is going to be 1 plus c over 2.

Now I've framed everything here, and let's just do a sanity check as we always do. If c increases, the price is increasing, and that should make sense. If the cost if something is more expensive to produce, the price of producing it-- the price that the seller charges is going to be higher.

Now this all seemed maybe-- we talked-- we used some fancy words, incomplete information, but it turns out, this is really equivalent to a very classical kind of Econ 101 problem, and I want to show the connection here.

So I know some of you have taken econ, maybe some of you haven't, but if we just look at the classic supply and demand curve, we have price on this axis and quantity on this axis. Well, we can think of a seller facing a buyer like this as facing a demand curve that looks like this. Because if I'm the seller, whatever price p I offer, we already said from our calculation over here that the buyer is going to buy with probability 1 minus p. So I'm effectively facing a demand curve that looks like this.

If I charge a price of 1, the buyer never buys, so demand is 0. If I set a price of 0, the buyer always buys, which means their probability of buying is just 1. That's my demand here. And if I set an intermediate price of p, well, then the buyer buys with probability 1 minus p, and it's this point along the demand curve.

So as usual, in economics, we see a downward-sloping demand curve. The lower the price I offer as the seller, the more likely the buyer is to buy. Now I'll highlight-- I'll say up front that normally we think of q as the number of units purchased, not the probability of purchase. But if we're facing a population of, say, 100 people, well, we can reinterpret this probability where a 1/2 probability of buying now means, in expectation, 50 out of 100 people in your town are going to buy. So we can reinterpret quantity as-- let's reinterpret it here as probability of buying times the number of people.

Up here, I just assumed there was a single buyer, but the math doesn't really change if you just multiply everything by 100-- say there's 100 people. And here, we see the connection with-- if you've taken econ and we talk about the monopolist pricing problem, this is exactly what they're doing.

Of course, I need one more thing on the figure, I need the cost, because that's also relevant to the monopolist. So let me put the cost c here. And what we can see is this box here represents the monopolist profit. Why? Well, the height of the box is p minus c. That's what we had up here for the profit per unit. The width of the box is 1 minus p. It's determined by the value of the demand curve at the price p that's offered, so that represents the probability of sale.

And therefore, what's the seller's goal? They want to move this price up and down and choose the value that makes this box have the largest possible area. And it turns out that this is maximized when the box is a square. And that's actually what happens here. When p star equals 1 plus c over 2, then-- let's just observe this. If p star equals 1 plus c over 2, then that means p star minus c equals 1 minus c over 2. So I have 1 plus c over 2. I'm subtracting c, but since I'm dividing by 2, I can subtract 2c on the top, and then I get 1 minus c over 2.

And then what is my demand? 1 minus-- oh, did I do something wrong here? Yes. We know what the answer has to be-- let's see. We'll write it out more carefully. Maybe we can write this as 2 over 2 minus 1 plus c over 2. And now it's clear. 2 minus 1 is 1, minus c, we get 1 minus c over 2.

So this is at p star. The height of the box is 1 minus c over 2, and the width of the box is also 1 minus c over 2, and we get a square. That's a special property of the uniform distribution. So one final comment is, when we drew the connection between the demand curve figure and the incomplete information figure, the distribution of the buyer's value exactly corresponds to the shape of the demand curve. So when we draw different demand curves in econ classes, one interpretation of that is it corresponds to different distributions of values in the population.

A demand curve shifting out means more people are willing to buy at a given price, which means the distribution of values in the population tends to be higher. OK, but this was all still a review. Let's now move into the incomplete information case-- or the more complicated incomplete information case. So let's go back over here.

So before, what we assumed is that-- well, one question here is to say, why does the buyer care? Why does the seller care about the buyer's valuation? They care because the buyer's valuation affects whether they buy. But it doesn't affect how costly it is for the seller to produce the item.

But there's a lot of cases where the buyer's private information-- so let's say-- I'll say this is another case. It's the case of interdependent-- sometimes called interdependent values in the notes. Interdependent values. And this is the case where the buyer's private information is relevant not just to how much the buyer values the good, but the cost to the seller of providing the good. So the buyer's private information is about v and c.

And this might be a bit strange. Why would the buyer know how costly it is for the seller to produce the item? But there's a few classic cases here. Does anyone know what's a classic case where the buyer might have private information that's relevant to both what the buyer gets out of the exchange and also the cost to the seller of making the exchange? Yeah?

AUDIENCE: Is it possible that the buyer would value the good even more and know that it's costs more [INAUDIBLE]?

IAN BALL: Yeah, that's-- I think that's possible, yes. Yeah, I think that's along the right track, yeah. There's one classical application that everyone talks about. So this is health care. Health insurance. And then I'll give one more example that's a bit closer to what you suggested. Health insurance. What does the buyer know that's relevant to both v and c in the context of health insurance? Yeah?

AUDIENCE: Their own status?

IAN BALL: Their own health status. The buyer is, here, an individual, they're buying insurance. This is one reason why we see the government quite involved in health insurance, and the reason why health insurance markets, unfettered private markets, tend not to work well, precisely because of this issue.

Because the buyer has private information that's relevant both to how much they want the insurance, but also to how costly it is for the seller to provide insurance to them. So what the buyer knows their health status. Of course, this is a gross simplification, but in short, you might think someone who's healthy has a lower v and a lower c. If I'm healthy, I don't care about insurance as much. And also, insurance is cheaper to provide to me because I rarely go to the doctor and the insurance company doesn't have to pay much.

On the other hand, someone who's quite unhealthy is going to value insurance a lot more. They're going to go to the doctor a lot and get a lot more out of the insurance, but that's also going to be more costly for the insurer because the insurer has to pay every time every time they go.

Now, health insurance markets sometimes try to mitigate this effect by offering different prices based on different observable characteristics. It may be that based on your price, older people tend to charge more-- charge more for health insurance than younger people, but there's still usually additional private information you have. Not all people who are 75 have exactly the same health, and they usually have private information about how healthy they are.

So this can create a problem, which is often known as adverse selection. And this term actually-- it's used in economics a lot, but it actually was first-- the term was first coined by people who worked in insurance and noted this phenomenon, and then economists subsequently studied it.

So let's look at quite, maybe, a stylized model of this. So let's represent the population or the buyer. The cost of providing insurance to the buyer is going to be uniform over 0, 1. So we can think of c as being a measure of the health of the buyer. I guess a higher c is someone who is less healthy because it's more costly to provide them insurance.

And then we want there to be some relationship between c and v. So let's assume that v is equal to alpha plus beta c where alpha is greater than or equal to 0 and beta is strictly greater than 0. So maybe it's helpful to draw a picture here to see what's going on.

We have c, we have v. c is uniformly distributed between 0 and 1, and v is alpha plus beta times c. So let's put alpha here and alpha plus beta here. And notice that when c equals 0, v equals alpha; and when c equals 1, v equals alpha plus beta. So we can think of people as lying along this line that looks like this.

We could have a more complicated model where maybe your cost doesn't directly pin down your value and vice versa, there might be some noise here, some individual heterogeneity. But here in this really simple model, everyone lives along this line. Here, we have the healthiest people who only get a value of alpha from the insurance, but only cost the insurer 0. And then as you move along the line, we have the least healthy people who cost one to the insurer to insure, but get utility alpha plus beta from the insurer.

In reality, we might need a two-dimensional model, maybe health and how rich you are, and maybe everything would be filled in and we'd have a distribution, but we're going to restrict just to this line to make things concrete. And because c is uniform, that means people-- we're imagining people are uniformly distributed along this line. Yes?

AUDIENCE: Just trying to understand qualitatively. How can someone did a lot of value out of the insurer, but not cost a lot or vice versa?

IAN BALL: It could be-- so wealth would be the easiest case. So someone who is really rich might be willing to pay a lot. So what is value-- yeah, it gets into what does value mean? Let's say I go to the doctor and fix my toothache. How much am I willing to pay to fix my toothache? I might be willing to pay less than someone else, even though it's fixing the toothache. So the cost to the insurer is just literally, how much does the doctor get paid, and the facility, to perform the operation?

But certainly, different people in the population would be willing to pay very different amounts for the same operation. A billionaire would tend to be willing to pay more than someone who doesn't have any money. So in this model, we're not taking into account that effect. We're looking at a simplified model where there's a one-to-one relationship between cost and value.

But yeah, I guess the way to think about value is value is about your individual willingness to pay, because remember, going back to here, this is almost the definition of value. If your value is v, it means you're willing to pay any amount up to v to get that item. And some people are willing to pay a lot more for something that has the same cost. Yes, question?

AUDIENCE: It's, like, not necessarily true that a plus-- that alpha plus beta is greater than 1. So are there any constraints that would make alpha plus beta greater than 1? Or in the real world, is it possible that alpha plus beta is less than 1?

IAN BALL: I see. Yeah, probably-- no, I think it could go either way. So I guess one example would be-- let's think of a-- there might be a very cheap surgery that really improves my quality of life. So in that case, the value that I get out of that surgery, my willingness to pay might be much, much higher than the cost of producing that surgery.

On the other hand, it might go the other way. There could be a really expensive new treatment that doesn't really help me out very much, and it may be hugely expensive to provide this treatment, but the utility gain and my willingness to pay for that treatment might actually be quite low.

So I think in general you could have either ordering between v and c depending on the context that you're in. I mean, this is a very stylized model. I'm not saying-- if we wanted to really directly conduct health care policy, we'd have to use a more complex model, but just to convey the ideas, I want to do that.

But yes, I think there could be-- and in fact, this is an issue, when people have full insurance and they don't have to pay to go to the doctor, they sometimes go to the doctor and get procedures that are extremely expensive to provide, but don't actually help them that much because they don't have to pay, and that's one kind of issue with insurance markets, yeah. Good. Any other questions on this?

OK, so now let's try to work through-- an analogous problem is over there. We still have the seller choosing to post a price, but now what's different? Now, it's only going to be the high-value people who buy at a given price. But what's the key difference here? What's going to be true about the people who choose to buy the insurance? If I set a price, only the people who value it more than that price are going to buy. What's going to be true about those people? Who's going to buy my insurance? Yeah?

AUDIENCE: The people that have high costs as well.

IAN BALL: The people that also have high costs. And this is the fundamental idea of adverse selection. If I go out into the market and I say this is the price I'm going to offer, then the people who are going to buy the insurance are going to be the people who are most expensive to insure, and that can create problems.

My favorite example of adverse selection is in the 1990s-- I don't know if I told you this-- I tell this story every year, sorry, I forget which year I told it. In the 1990s, United or Delta, one of the airlines, offered an unlimited flying program where you could fly forever. And it cost, I think, like $200,000 in today's terms. And Delta thought, well, most people don't fly that much. $200,000, that's a lot. It was a disastrous program for the airline because there's some really weird people who just like to fly.

And one guy just took the program and just flew first class around the world every single day for the rest of his life. That's an example of adverse selection. When you post a price, the person who buys it is not the average flyer, it's the person who really values flying a lot. And this was a person who was extremely costly to give an unlimited flying license to.

For me, probably 200K, the airline would make money on me. Of course, my behavior would change once I got it. But in this guy's case, they lost a lot of money, and they ended up going to court and trying to nullify the contract, and it was this big, big mess.

OK. I guess another adverse selection effect might be health insurance. Another might be in financial markets. If I'm selling a mortgage-backed security and I don't know much about the value of this mortgage-backed security and someone's willing to buy it from me, well, the fact that they're willing to buy it probably means it's very valuable and I'm not getting a good price. If they're not willing to buy it, then it goes the other way.

So there's always the problem, you always have to think, why is this person willing to buy from me? And often, if they're willing to buy from me, it might mean I don't want to sell it to them, and this is the fundamental issue of adverse selection. If someone's willing to buy an unlimited flight ticket for 200K, maybe it's not a good idea to sell it to them.

OK, but let's now get into the math here. So let's suppose I offer a price of p. So the seller offers a price p, then the buyer buys if and only if-- well, as before, their value is greater than or equal to p. And now, instead of saying p is between 0 and 1, let's say p is between alpha and alpha plus beta. Because, again, if I charge a price below alpha, no one's going to buy, and if I charge it above alpha plus beta-- sorry, below alpha, everyone's going to buy, and above alpha plus beta, no one's going to buy, so we can just focus on this range-- this range of value space.

So c buys-- or the buyer buys if and only if v is greater than or equal to p. Well, let's write that out. That means alpha plus beta c is greater than or equal to p. Here, I've just written out what v means. And now let's solve for c, and we get c is greater than or equal to p minus alpha over B.

So now we can compute-- let's actually go over to this board here. Let's compute pi of p, which is what is the seller's expected profit if they post a price of p? Now over here, we just wrote p minus c times the probability that v is greater than or equal to p. This part stays the same. This is the probability they buy, but now we have to adjust p minus c.

Because now, there's not one c, there's a population of people. And the c that we care about is not the average sea in the population, but it's the average c among the people who buy. So we crucially need a conditional expectation.

Maybe I'll write it this way. p minus e of c, given that someone buys, and that means given that c is greater than or equal to p minus alpha over B, over beta. And then I'm going to multiply that by the probability that v is greater than or equal to p-- or maybe I'll write it the same way. I'll write this as c just to make the math a bit easier.

And now we can see this pretty clearly in our graph. Let's say I post a price of p here. Well, who's going to buy? It's going to be people along this segment of society. These are going to be the sickest people, the people who value insurance the most. And a naive mistake that I think historically some companies made is they did this calculation just taking the expected cost in the population.

Not recognizing that, well, I shouldn't take the expectation over the entire line, I need to only take the expectation among the people who are actually buying my insurance. Or in the airline example, only among the people-- the flying frequency of people who actually buy an unlimited flying ticket.

So now let's try to compute this expectation. Well, it turns out, c is uniformly distributed. Down here, this is actually p minus alpha over B, over beta. So I'm only focusing-- I'm only actually selling to people who lie in this range. And because c is uniformly distributed, the average cost among people in this range is just going to be the midpoint of this range.

So this is going to be p minus-- well, the midpoint of this is going to be 1/2 times 1 plus p minus alpha over beta. So it's still price minus roughly cost, but now it's expected cost conditional on buying, and then I'm going to need to put this probability. Well, this probability is just the length-- or the length of this segment, which is 1 minus p minus alpha over beta.

And let's-- I think it's-- we're interested in p, so let's try to do a little algebra to get p by itself here. So in this, we have p here, and we have minus 1/2 times 1 over beta. So we get p times 1 minus 1 over 2 beta minus 1/2 times alpha over beta.

So now let's analyze what's going on here. This is the purchase probability. And this is the profit per purchase. Now, as usual, the purchase probability is decreasing in the price. This is the effect we always see. When I'm a firm and I charge a higher price, fewer people buy. That's very intuitive.

The new effect is here. This is saying, what profit do I make per person who buys from me? And notice that the coefficient on p before, the coefficient on p was 1 When I charged a price, when I increased the price by $1, every time someone bought from me, the amount of revenue I generated-- the profit that I got from that person increased by $1.

But now, it doesn't increase by $1. It increases by $1 minus 1 over 2 beta. So let's understand this. I raise my price by $1. I don't make a more per person I'm selling to, I only make one minus 1 over 2 beta more, which is strictly smaller. I don't get the full dollar. Why is that? So why is this less than 1? Yeah?

AUDIENCE: Because as you raise the price, the people who will be buying it get more value out of it, so they'll cost more [INAUDIBLE].

IAN BALL: Exactly right. So this is reflecting two effects. The first effect is, yes, when I raise my price by $1, my revenue per patient increases by $1. That always happens. But the second effect is when I raise my price by $1, I change the composition of the people who buy from me. I shift the people who buy from me to a less healthy subset of people, and on average, those people are costlier to serve. And in fact, this is showing they're costlier by exactly 1 over 2 beta.

Where does that come from? Well, we can see here, these are the segment of people I serve. When I raise my price by $1, I increase the left endpoint of this segment by 1 over beta, and therefore, I increase the average of this segment by 1 over 2 beta. So this is exactly the-- 1 over 2 beta is the marginal effect of my price on the average cost among the people who buy from me. More precisely, it's the derivative of this expression with respect to p.

And we can see-- let's look at different cases of beta. So what if beta is large and what if beta is small? So if beta is very large, adverse selection is not much of a problem. If beta is very large, this curve is very, very steep. And if beta is large, it says that there's not a huge effect on the average cost when I change the price.

Because if beta is large, it says that c doesn't change much with v. Another way-- I think the easiest way to see this is to solve this expression for c as we've done. And so note, c equals v minus alpha over B.

So when I change the price by a little bit, I'm changing by a little bit the value of people who are buying from me. And the question is, if someone's value changes a bit, someone's value of buying for me changes by a bit, how much does my cost of providing to them change? And we see, it's exactly proportional to 1 over beta.

So when beta is really, really large, adverse selection is not much of a concern. So little adverse selection. When beta grows small, then people who value the item just a little bit more are extremely costly to insure.

So we have severe adverse selection. And let's look at the case where beta is greater than 1/2. Or, sorry, less than 1/2. So this is the extreme case. What happens if beta is less than 1/2? What happens to this coefficient? Yeah?

AUDIENCE: It's negative?

IAN BALL: It's negative. So what does that mean? I raise prices and I actually earn less per person than I did with the lower price. What it says is if beta is less than 1/2, the shift in the composition of people who buy from me more than outweighs the fact that I'm charging a higher price. And I'm actually earning less per patient I serve. And what that means-- maybe we'll call this market breakdown.

Why? Well, if beta is less than 1/2 then this is decreasing in p, and this is also decreasing in p. So that shows my profit is certainly decreasing in p over the range alpha to alpha plus beta. So when I say market breakdown, I guess the only possibility is-- maybe what I mean is screening is not possible. Because the optimal price is either p equals alpha or no sale.

If my profit is decreasing in p, the only possible maximum of this expression would be at p equals alpha at the left endpoint. So I either have to serve everyone in the market or serve no one. And it may be optimal actually for me to serve no one. But what it says is there can't be a screening effect. I can't charge an intermediate price and get some people to buy and some people not to buy. The only possibility is I set the price low enough that everyone buys from me or I just don't sell at all.

Maybe we can interpret everyone buying from me as kind of universal coverage. We can't have a premium insurance plan that only certain people buy. That can never be optimal if adverse selection is severe enough. Yes?

AUDIENCE: And so if you sell to everyone at a price of alpha, is your profit supposed to be 0? Is that what the [INAUDIBLE]?

IAN BALL: It's going to depend on the values of alpha. So in some cases-- so if you set the price to alpha, you're going to get some number. If that number is-- if your profit is greater than 0, that's the optimum. If it's less than 0, you're better off selling to no one.

AUDIENCE: Oh, so your point is just that if beta is less than 1/2, then the only option-- like, you have to solve--

IAN BALL: Yes. There's no interior solution. Either I sell to everyone or I sell to no one. I mean, we could go more into details depending on the details of alpha and beta, but basically, we have a function, and it's going like this. So its maximum is at alpha equals 0, yeah.

We have to be careful if sometimes things can be negative. So, yeah, when I say it's-- OK, we have to be careful. When we say it's decreasing in beta, that's not true if one of these expressions is negative, so let me make the more precise statement that the solution occurs either at alpha or beta-- at alpha plus beta.

AUDIENCE: And so why does it make sense that you would sell to everyone? Like, what about beta makes it so that you could potentially make money from selling everyone?

IAN BALL: So whether it's profitable to sell to everyone is really going to depend on alpha. What beta is controlling is what happens when I slightly raise my price. So when beta is very, very small, it says adverse selection is so severe that I would never want to sell to only some people in the population because by raising my price a little bit, I'm actually earning less per person because those people are so much sicker.

So when that effect is really strong, I don't want to be vulnerable to adverse selection. I don't want to sell to the subset of the population who value it most. I either want to sell to everyone or sell to no one. And that's going to be controlled by alpha. Alpha could be enormous.

If alpha is a million and c is pretty small, then I certainly want to sell to everyone. Alpha is kind of about a level effect. It's like, is this a good market to be in? Do I want to sell to people or not? Beta is a marginal effect, which is controlling what happens to my profit as I slightly change my price. And it's saying, if beta is small enough, the adverse selection effect is so strong that I never want to slightly shrink the set of people I sell to because the cost among those people is going to be so much higher. Yes, again.

AUDIENCE: I'm thinking about real-world examples where you have insurance markets, I'm thinking like disaster insurance where if you live in a certain area, you have to buy insurance, or if you have health care given by your employer. Would that be a case where you just end up charging alpha because of the policy intervention? Like the insurance company charges alpha for everyone.

IAN BALL: Yeah, so I guess a separate issue that comes up with disaster insurance is that often, disasters are correlated. So you have a heart attack, someone else is more likely to have a heart attack. If you get hit by a hurricane, everyone's often hit by a hurricane. So often there's an issue that these disasters are too big to insure. So often, the state steps in and provides the insurance because no private company can do it, and then the state puts in these regulations.

I think really, what's going on in those cases-- there's a lot going on with disaster insurance. In this case, here, I think normally, with reasonable parameter values-- I mean, this is a very simplified model. With reasonable parameter values, when adverse selection is bad, the solution is probably going to be to sell to no one because you're going to have-- if you have any really healthy people who are willing to pay a very low amount for insurance--

I mean, think about it. Let's say-- if I'm supposed to insure everyone in the population, but I can only charge a price that makes the healthiest people willing to buy, then that's probably not going to be profitable for me. So I think what I would say is, in realistic settings, the alpha where markets break down, the alpha is going to be small enough so that there's not going to be a sustainable private insurance market.

Of course, there's a separate issue here that we haven't-- there's a separate question, which is how well can we monitor people? So another thing that could happen is you might impose very strict health standards. You might say-- there might be a functioning insurance market, but only for people under 50.

And then people over 50 are just excluded from the market. And that conditioning is not present, is not captured here, or at least the interpretation of this would have to be after we conduct medical tests and after we condition, and then beta might become bigger, yeah.

Yeah, I think the-- yeah I think the insurance market, the disaster insurance is interesting, but I think it's really a separate issue because people don't really have private information about how likely their home is to flood. I mean, maybe if the government has to use really blanket policies and they're not fine-grained enough, but at least in the private market, I think that's not really an adverse selection issue, it's really a too-big-to-fail issue. Yeah, but it's an interesting question.

And, yeah, actually one of my colleagues is just writing a paper about this in California wildfire insurance and looking at how insurers handle this. Yeah, it's interesting.

OK. We've talked a lot about screening. Let's now talk about signaling. And this gets a little messy, so we're going to have to work with an even simpler model. And, yeah, we should have time just to finish this off.

So now let's finally move to the case of signaling. So we're still with a seller and a buyer, we're still in incomplete information bargaining, but now it's the seller who has private information. And now we're going to say the seller has private info about v and c.

So this is another canonical application where the seller-- we said insurance is the canonical case where the buyer has private information that's relevant to both v and c. Can you think of a setting where the seller has private information about v and c? It's kind of a classic story here.

Used cars, kind of the classic case, or more generally, used goods. Here, I think it's important to think of c maybe not so much as a cost, but more as how much the seller values the good. And the idea is if I'm going to buy a used car, the seller of that used car knows whether it's making a bad sound, whether it's had a lot of accidents, whether there's these problems. Basically what they know is the quality of the good.

And that quality, unobserved quality, is relevant both to how much the buyer values it and how much the seller values it. But because the seller has had the good for a while, they know a lot more about it than you as the buyer who come by and just try to inspect your good.

And the classic case would be used cars. So there's a very famous paper called "The Market for Lemons" by Akerlof. He won the Nobel Prize for identifying this issue of why used car markets often don't work well, and it's the same classic thing that we said, going back to the defendant and the prosecutor, if someone comes to you and says, "I can't sell my used car, but I'm going to give you a really good deal. It's a really good deal, it's really cheap." What are you going to say? "Wait, why is it so cheap?" It's cheap because it must be a lemon. It must be-- that's another word for a car that doesn't really work. And for that reason, these markets can break down.

I think my financial example was another good one. So mortgage-backed securities might be another case where it's not clear which side has more private information, it depends who they are, but in some cases, the holder of the mortgage-backed security may have private information that's relevant to the value of this asset, and both the buyer and the seller care about that.

OK, so let's look at a really simplified model where we're going to say that the good has quality either L or H. And here, these are not numbers, these are just labels. It's either low quality or high quality. But the quality the good is relevant to both the buyer and the seller. So if it's a low-quality good, then-- and, I don't know. All these inequalities may be confusing.

So if it's low quality, then the seller's cost is cL and the buyer's value is vL. If it's high quality, the seller's cost is cH and the buyer's value is vH, and we have all these inequalities. For the seller, a low-quality item is less costly than a high-quality item, but we also want the seller's cost to be lower than the buyer's value for a low-quality item.

For a high-quality item, the seller's cost is smaller than the buyer's value, so there are potential gains from trade, but again, the buyer values it more when it's high quality than when it's low quality. So there's a string of inequalities here.

And now what we're going to look at is price signaling. I'm going to set a price to the seller, the buyer is going to see that price and infer something about the value of the item, and then we want to try to solve for an equilibrium of this. And we're going to have to use perfect Bayesian equilibrium here to analyze this.

Now we need to specify a prior distribution about the quality of the good. So let me say a little bit more here. The seller knows quality. The buyer doesn't know the quality, but they just know the distribution of quality in the population. And we'll say that-- we'll put pi here and 1 minus pi here.

So the prior distribution is that fraction pi of things are high quality, fraction 1 minus pi of things are low quality, the seller actually knows the quality, the buyer does not.

And after observing the price, the seller is going to set a price. The buyer sees the price and is going to update their beliefs about the quality of the item based upon the price.

So let's try to construct-- and I think you'll have more practice doing these in the problem set, but we'll start constructing PBE in this game. Perfect Bayesian Equilibria. And as usual, we're going to separate into two cases. We're going to look for a pooling equilibrium and a separating equilibrium.

So on the one hand, we'll do pooling, and the other is separating. A pooling equilibrium means sellers with high-quality products and sellers with low-quality products both sell at the same price. So the price alone doesn't give away whether it's high quality or low quality. So in a pooling equilibrium, the seller's strategy-- so let's make sure we understand what's going on here. What is the seller's strategy in this game? We need to say if the quality is low, what price do they set? And if the quality is high, what price do they set?

In a pooling equilibrium, these are going to be the same price, and we'll call that price p star. And this is going to be a parameter of the equilibrium that we're going to try to solve for. What is the buyer's strategy? As I say, I think writing down the strategies correctly is often the hardest part of this. So what is the strategy for the buyer in this game?

AUDIENCE: Can't you find an expected value of the quality in some way? Or--

IAN BALL: You could-- you could, but let's not even think about optimality. Let's just think what kind of object is a strategy? It's a number, is it a vector, is it a function? Just what thing is it? What kind of mathematical object? I think that's often the first step. Yeah, in the back.

AUDIENCE: Would it be either accept or reject as a function of the price?

IAN BALL: Exactly. As a function of the price of the seller offers. Exactly. So it's going to be a function-- maybe I'll call it q since we think of q as quantity of buying. q from maybe-- if we don't restrict the price at all, well, the seller can offer any non-negative price, say between 0 and infinity, not inclusive, to-- well, maybe we'll call 0, 1.

So q, we can think of this as the quantity that they buy. They either buy 0, that means they don't buy or reject, or q equals 1, that means they buy or accept the offer. And let's put our graph here. We have vL, we have vH. And I want to try to work up gradually and try to build up an equilibrium.

So the first, I think, useful observation when we're trying to pin down the equilibrium is to say, well, there are some prices where we know what the buyer has to do, whatever their belief is. So let's think.

The buyer's belief is going to form a belief about the quality of the item. And if they're really pessimistic, they would think the item is worth vL. If they're really optimistic, they're going to think the item is worth vH. And if their beliefs vary, they're going to value the item somewhere between vL and vH depending on how likely they think it is that it's high quality or low quality.

But if they get offered a price below vL, what are they going to do? They have to buy because even if they're certain the good is low quality, this price is still lower than their value for the good. What if they see a price above vH? They're never going to buy it. Doesn't matter-- even if they're certain the good is high quality, it's not worth buying because vH is smaller than the price.

So we immediately know that-- let's think of q over here. We're immediately going to have 1 here and 0 here. And the flexibility is going to be between vL and vH. For a price between vL and vH, whether they buy is going to depend on their belief, because if they're certain it's high quality-- let's say there's a price in here. If they're certain it's high quality, they do want to buy, if they're certain it's low quality, they don't want to buy, and it's going to depend on what their belief is.

OK. What else can we say when we're looking for an equilibrium? Well, remember, PBE, we have to have sequential optimality for the seller, sequential optimality for the buyer, and then Bayes' consistency for the buyer. Now Bayes' consistency only pins down the buyer's beliefs when they see a price that they're supposed to see in equilibrium. It only pins things down on-path. So what is going to be on-path in this equilibrium? What price are they expecting to see in this equilibrium?

Just p star. That's the only price they should see. Because that's what the low type seller is offering and that's what the high type seller is offering. So if they see a price of p star-- so let's-- what belief do they form? They apply Bayes' rule. What happens when they see a price of p star? Yeah?

AUDIENCE: We would have been given pi, we would have an intuition about is it low type or high type, and then-- like, we would know which one of those is higher probability mitigating whether the car is actually low type or high type.

IAN BALL: Right, but here, we're looking for a pooling equilibrium. Where both the high-type seller and-- both the low-type seller and the high-type seller sell it at p star. So when I see the price of p star, what is my belief? Yeah?

AUDIENCE: That it's high type of probability pi.

IAN BALL: Right. It's just the same as the prior. So if everyone is charging a price of p star, I keep my beliefs. My beliefs stay exactly what they were at the prior. I have no reason to update my belief-- I mean, formally I'm applying Bayes' rule, but when I update my beliefs, they're the same as they were before because everyone offers the price of p star. The price doesn't really reveal anything. So maybe I'll say p star belief pi on H.

Now, if I believe the item is high with probability pi, then the value I assign to it-- maybe I'll call the value v star-- is going to be pi times vH plus 1 minus pi times vL. That's my expected valuation of the item given that my belief here is pi. So we're kind of building this up. We're looking for an equilibrium where the seller charges a price of p star.

We know that the buyer's valuation has to be this. And we want the seller to sell the item at a price of p star. So what must be true about p star relative to v star in this equilibrium? Yeah?

AUDIENCE: p star equals v star?

IAN BALL: It could equal v star, but it has to be at least weakly smaller than v star. You might-- it could be v star. So let's put v star in here. And p star is going to have to be somewhere smaller than v star. And we know that at a price of p star, let's look for an equilibrium where the buyer does buy at a price of p star.

Now the trick for these equilibria is often you want to make deviations as unattractive as possible. So we want to make this an equilibrium. That means we want to make sure the seller doesn't want to offer a price different from p star.

To make sure the seller doesn't want to offer a price different than p star, it better be the buyer never buys if a price is strictly higher than p star. So if I see a price higher than p star as a buyer, I'm not going to buy, and that way, the seller won't want to deviate to this.

So let's look for an equilibrium that goes like this. Seems like a natural equilibrium. If I see any price that's weakly below p star as a buyer, I buy. And if I see any price above p star, I don't buy.

Now we have to specify beliefs. I think the easiest thing is to say-- well, let's say-- it doesn't even-- it doesn't even matter, but sometimes this is a good trick. Let's say that if I see any price strictly below p star, here, I have a lot of flexibility in what the beliefs are because we're off-path, but I know that the buyer is buying, so to make that optimal, let's just say they believe the good is certain to be high quality. It doesn't really matter what I do here.

If they see a price strictly greater than p star, they believe it's low quality and don't buy. And if they see a value-- a price of exactly p star, well, we know what they have to believe. They have to believe the good is high quality with probability pi because this price is on-path, and therefore, Bayes' rule can be applied.

For these lower prices and these higher prices, I'm totally flexible in how I choose the buyer's beliefs because these are off-path, and if I try to apply Bayes' rule, I would divide by 0, and therefore, Bayes' rule doesn't have any bite.

OK, so I think we're good, but I think we need one more constraint. So we've defined p star. We've said p star is between vL and v star. This means that the buyer is willing to buy. I think there's one more potential issue with this equilibrium, so let's go through it.

We've checked-- maybe I'll just write this here. We've checked belief consistency because here, the buyer believes it's low quality. They believe it's high quality, they certainly want to buy. Here, they believe it's low quality, they certainly don't want to buy. Here, they believe its expected quality is v star and they're, therefore, willing to buy at a price of p star. So Bayes' consistency is right.

Receiver optimality, RO-- ah, receiver-- buyer optimality is good, but I think there's one more issue with seller optimality. So we have to check, does the seller ever want to deviate to offering a different price?

Well, the trick is always to say, what is the most attractive deviation? Would the seller ever want to charge a price-- a lower price here? No, because they're already selling it with certainty. If they offered a lower price, they'd still style it with certainty. But what might the seller want to do? What do we have to check? Well, does the seller actually want to sell it at price p star?

Now the issue is, there's two types of sellers. There's the high-type seller and the low-type seller. We have to say, what is the more tempting deviation? The issue we have to worry about is-- and maybe I'll do separating on another page. Let's look at both cases. The high-type seller is comparing p star minus cH. These are the two constraints.

If I'm a high-type seller, I could charge a price of p star, and I sell at p star, but I have to pay, but my cost is cH. Or I could set a price that no one will buy it, and therefore, I get 0.

Conversely, if I'm the low-type seller, I could set a price of p star, pay my cost, or I could sell at a really high price that no one's going to buy it. Which of these is harder to satisfy? The top one, because the high-type seller finds it more costly to produce. So they might say this price p star, it's not really worth it for me. So this is going to be the binding constraint. This is the critical-- the important constraint. And this is going to be the final constraint that we need to impose.

So what we need, to write it in words, is we need vL and cH less than or equal to p star less than or equal to v star.

This inequality says the buyer is actually willing to buy at a price of p star. This inequality says even when the seller knows the good is high quality, they're still willing to sell at a price of p star. And the vL inequality is important because if p star were less than vL, then the seller could always slightly raise the price and the buyer would still always have to buy because we know the buyer always buys when the price is below vL.

So these are the three inequalities that we get. And I think when you go through and try to solve for PBE, these are the steps you have to go through. Let me say just a little bit about the separating equilibrium in the last two minutes. So that was pooling. Let's go down here to separating.

So for the separating equilibrium, we now have pL not equal to pH. So when the seller knows the good is low quality, they sell it at pL. When they know it's high quality, they sell it at pH. These are different. As before, the buyer is choosing a strategy q from 0, infinity to 0, 1.

Now, which prices are on-path? It's both pL and pH. So now let's look at beliefs. If the buyer sees a price of pL, they know that the seller is low quality, so they assign probability 0 to h. If they see pH, they know the seller is high quality, so they assign probability 1 to h. And if they see any other price, well, now their beliefs are flexible because we can't apply Bayes' rule, this is off-path, so flexible.

And we can go through and try to do the same analysis. Maybe we'll just say two kind of tricks here. You might look-- let me just say one final thing. You might try to find an equilibrium where the buyer buys at price pL and at pH. I'm going to argue that can't be in equilibrium. Why could there not be an equilibrium where the buyer buys when they see pL and buys when they see pH? Why can't there be a separating equilibrium where the buyer always buys?

Think about the seller's incentives. There are two prices at which the seller always buys. What would you do as the seller? Well, one of these is higher. Let's say pH is higher. If you were the seller, you'd always sell at pH because you'd get more revenue and the buyer is willing to buy.

So it turns out, there's not going to be a separating equilibrium where the buyer always buys. We're going to actually have to have a mixed strategy where the buyer buys with some probability following pH, but you'll see how to work through that in the problem set. But let's make sure we understand the intuition for that. OK, let me stop there. I'll see you on Thursday.