[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: So today we're going to consider an application of backward induction. Today is going to be about negotiation. And the basic idea is that when two parties are bargaining or negotiating, they often-- one party might make an offer and say, let's do this. And then the other party might respond with a counteroffer, and they might go back and forth.

So right now the government just shut down. That's exactly what's happening between Trump and the Democratic Party. They might be countering offers and go back and forth and see if they can eventually reach an agreement.

So the first model-- we're going to discuss two models of negotiation. One is going to be about pretrial settlements. And the other is going to be about price haggling. And I think these are kind of illustrative applications. Of course, we're not going to include all the details of these applications. We're going to introduce foundational stylized models. But I think they can capture some important elements of these issues.

So let's start with settlements. So in many cases, in criminal cases, we often have pretrial plea bargain bargaining. And in civil cases, we often have the two parties working together before the case and trying to reach a settlement to avoid trial.

So today, let's think about civil cases. I think that fits the model best. And in practice, a lot of civil cases never make it to trial. The two parties reach some settlement before the case goes to trial. What are some reasons why cases like this often don't go to trial? Why do we see so many pretrial settlements? Any thoughts? Yeah?

STUDENT: Both sides want to avoid having legal fees.

IAN BALL: Legal fees. So legal costs are a big part of it. And that's really what our model today is going to focus on, that having a drawn-out trial is very costly. Did you have the same idea or any other-- same idea. Anyone else? Any other reasons?

It could be legal costs. It could just be the cost of delay, which is similar but slightly different. It could be that even aside from paying the lawyers, if we're going to be stuck in court for so long, and it's going to delay our productive activity, that might be costly.

Another big one could be risk aversion. This is not going to be in our model today. But it might be that I think that if we go to court, there's a chance I win tons of money. There's a chance I don't win anything. So I'd rather accept a settlement where I win some small intermediate amount of money with certainty. So risk aversion, and I should be clear here, over monetary payoffs.

I think this is also true in criminal cases where someone might think, well, there's a chance I get a life sentence. I'd rather accept a plea bargain where I'm certain to go to prison for five years, rather than have a risk that I serve a life sentence.

And then maybe one final one that we often see is about privacy and maybe embarrassing information coming out, or potentially trade secrets being revealed at trial. There are some provisions to try to prevent this, but I think there's still some risk of this.

So today, we're going to introduce a formal model that can try to make some predictions about what kind of settlements we'll see and how long negotiations will last. Of course, this model isn't going to capture everything, but it's going to be a good starting point for today.

So in our model, we're going to have two parties who we call the plaintiff and the defendant. So the plaintiff is the party that's suing someone else. So the plaintiff feels that they've been wronged, and they're seeking a monetary settlement.

The defendant has been accused of wronging someone, and they're being potentially asked to pay a monetary settlement. So let's be clear. At the end, there's going to be-- if they reach some settlement, that's an amount of money that's going to go from the defendant to the plaintiff. So from-- maybe I'll just write it with arrows-- D to P.

So just to make sure we're on the same page, the defendant is hoping to pay out a very small settlement. The defendant doesn't want to have to pay a big fine or a big punishment to the plaintiff. The plaintiff is seeking a very large settlement because they want to get paid out as much as possible.

And in this case, we're going to make a few simplifying assumptions. One of them is that there's no uncertainty about the outcome of a trial. Why are we doing this? Well, we want to isolate the effect of legal costs.

In this model, we're not going to worry about risk aversion and the potential uncertainties that come with trial. So what we're going to assume here, to make things simple, is that if they go to trial, the settlement amount is known to be-- will definitely be J. And we'll use to indicate judge. This is what the judge decides. So both parties know if we go to trial the amount that's going to be agreed upon.

So here, we're not thinking, you might imagine the judge deciding whether the plaintiff-- whether the defendant did anything wrong at all. Here, we're imagining that yeah, the defendant has certainly done something wrong. The question is only how much they have to pay. And here, the assumption is that if they go to court, the amount of the settlement will be some fixed amount, J.

And then we need to specify what the cost of going to court are. So maybe the trial costs are going to be Cp for the plaintiff and Cd for the defendant. And I think in the notes they use capital letters. But I find at the board it's hard to know what's capital and lowercase, so I'll just use a hat to indicate the costs if we go to trial.

But then each period that we negotiate before the trial also has costs. You don't just pay the lawyers when you go to trial. If you want to try to reach a settlement, that, in itself, is costly.

So we're also going to assume that there are per period-- and that will be clear when I describe the periods-- per period negotiation or pretrial negotiation costs. And these are just going to be cp and cd without hats. So each period that we negotiate for, the plaintiff has to pay a cost of cp to their lawyers, and the defendant pays a cost, cd, to their lawyers.

So now let's describe the alternating offer structure of this game. This is a very common format that we see in reality. You make some offer. The other side will either say it's acceptable, or they might reject your offer and then come back with a counteroffer that's more favorable to them.

So we're going to imagine that we have time going-- we have 0, period 1, period 2, period 3, all the way up to, let's say, 2n minus 1, 2n minus 1, and then 2n. And what we'll imagine is that it's known that the date of the trial is set. So this is the trial.

So we know that the trial is going to occur in period 2n. And if we go to trial-- when I say it will occur, what I mean is if we don't reach a settlement before then, then the trial occurs, and we know what's going to happen. The settlement is going to be J, and we're going to have to pay all these legal costs.

But in anticipation of the trial, what we're imagining is there's time to have n rounds, or 2n minus 2n periods, but n total rounds of alternating offers. So what we imagine is in the first period, the plaintiff is going to be the one who's going to propose.

So it's a little confusing. P is for plaintiff. You're not proposer. But the plaintiff, maybe I'll say offers, I'll call this S0. So in period 0, they're going to say, look, let's just settle. And the amount that you're going to have to pay me-- I'm the plaintiff, you're the defendant. The amount that you're going to have to pay me is S0. Do you want to accept this or not?

And if it's accepted, then you pay me this amount. We're done. And if it's rejected, we continue. So then the defendant can choose either to accept or reject. And if the defendant accepts this offer, well, now we're done. The defendant pays S0. We never go to court. The game is over.

But if they reject this offer, well, they're not just going to reject. They're going to come back tomorrow and propose a counteroffer. So in the next period, the defendant is going to make an offer. And now I'll call that offer-- they can make any offer they want-- S1.

And this is going to say, let's settle now, and I'll pay you S1 dollars if you accept it. And again now it's the plaintiff's turn who's going to say-- who's going to choose either to accept or reject. And we can see how this is going to keep going until-- let's see.

In even periods, the plaintiff is making the offers. So 2n minus 2 is an even number. So that means here, the plaintiff, if we get to this point, would make an offer, let's say, s2n minus 2. And again, in this case, the defendant is going to choose to accept or reject. If it's rejected, we move here.

Now it's the defendant's turn. They're going to make this last offer-- S2n minus 1, and the plaintiff is going to choose whether to accept or reject that offer. And this is only00 and only in this last period are things different.

So in this last period, if the offer is accepted, great. It's over as before. But if it's rejected, well, now the trial is here. So instead of going to one more round of negotiation, we just go to trial, and we see what the outcomes are going to be.

So just so we understand the costs, the Cp hat and Cd hat cost, the trial costs, are only realized if we actually go to trial. Whereas, the per period negotiation costs are realized in every period until we reach a settlement.

So here, we have Cp Cd. Here, we have Cp Cd and so on. So just so we understand, if we were to settle in this period 0, we would have only had to negotiate for exactly one period. So we'd each bear these costs only once.

But if we settled in, say, the third period, period three, well, it's tricky because we started at 0. So if we settled here, then we've had to negotiate 1, 2, 3, 4 periods. And that means we'll each have to pay four times our per period negotiation costs.

If we were to all the way go to trial, we're going to have to pay our trial costs, and we've negotiated each of these periods. So we'll each have to pay 2n for the 2n negotiation periods times our per period negotiation costs, plus the cost of going to trial.

So this is the game. And now what we'd like to do is analyze this game using backward induction. Any questions on the setup of the game? So let's first just recall what backward induction is going to deliver for us. So we're going to apply backward induction. What this is going to give us is a special Nash equilibrium.

Remember, first we have to think, are we allowed to apply backward induction? Well, backward induction only can be applied to games with extensive form games with perfect information. And here, we have perfect information because there's no point in the negotiation where someone doesn't know what someone else did previously. Everything is observed in this game. So we're allowed to apply backward induction.

The problem with games like this is that there are many, many Nash equilibria that may involve non-credible threats. So backward induction is going to give us this very special Nash equilibrium. And sometimes the term we use-- maybe another way of saying special is it's going to be a Nash equilibrium, but it's also going to have this property that's often called sequential optimality.

Remember, our key observation from last class was that Nash equilibrium by itself doesn't necessarily imply sequential optimality. That is, it may be under some Nash equilibria that players are behaving suboptimally at contingencies that are not reached. And sequential optimality imposes the stronger condition that we have optimal play at every contingency. Particular, even those contingencies that are not reached.

So Nash equilibrium is already going to take care of the contingencies that we do reach. Nash equilibrium is already going to require optimal play at those contingencies. But Nash equilibrium doesn't actually impose many constraints on these unreached contingencies. And what sequential optimality does, or requires, is that players are playing optimal even there.

Because if players were not playing optimal at some contingency, that would capture what we think of as a noncredible threat. That would mean a player was threatening to behave in such a way that would not be optimal for them or credible for them to follow through on if that contingency were actually reached.

Now, a key thing to keep in mind here is that what backward induction gives us is a Nash equilibrium. So it's going to give us a strategy profile. So often, when people analyze a game like this, it's tempting to say, oh the answer is we settle in period 5, and this is the amount of payments. That's an outcome of the game, but that's not a complete description of the strategies.

So a Nash equilibrium is going to have to specify the player's complete contingent plans throughout the game. And that means that the strategies in games like this are much, much more complicated than the actual outcome that's reached, which might just be we settle for this amount in this period.

OK, so let's get into it. It's backward induction, so we're going to work backwards, not surprisingly. And let's start with-- well, you might say we're working backwards. Let's start in period 2n minus 1. But actually, even within period 2n minus 1, we have sequential moves.

First, the defendant makes an offer, and then the proposer responds. So if we want to work backwards, we're actually going to split this last period up into two, and we're going to work in the second part of the last period.

So we're going to say suppose we're in period 2n minus 1. And the contingency we're at as the plaintiff, is that we're facing some offer from the defendant. So given some offer, S2n minus 1-- this is an offer by the defendant-- the proposer has two choices. They can either accept the offer, or they can reject the offer.

I should point out here that I'm being a little sloppy. I'm saying suppose the proposer receives an offer of S2n minus 1. But technically, this doesn't specify the full contingency. What would I need to specify if I wanted to say-- if I were to write out this whole game tree, and I want to figure out where I am in the game tree, what is this contingency going to specify within the tree? It's not just enough to specify the offer in the last period. Yeah?

STUDENT: Maybe that you need to know everyone before?

IAN BALL: Everything that happened before. Because we know in a tree, when we're at a node in the tree, that node tells us the entire path through the tree. So technically, a contingency would be something like, the plaintiff made this offer in the zeroth period, and it was rejected. The defendant made this offer in the next period. It was rejected. The plaintiff made this offer, and it was rejected, and so on.

But what I implicitly mean here is there's a lot of different contingencies, but the incentives of the plaintiff are going to be the same at all the contingencies where this is the offer that's made. So I'm basically grouping all the contingencies in my head and reasoning about all of them together at the same time.

Because at this point, if you're offering me S2n minus 1, I don't really care what you offered me three periods ago. That's not on the table. So it's not relevant to my optimal decision. But if I wanted to be really precise, I could, in principle, behave differently at different contingencies.

So now let's carefully write down my payoffs. If I accept this offer as the plaintiff, what are my payoffs going to be, my total payoffs, taking into account the settlement, my legal costs, all these things? Yeah?

STUDENT: [INAUDIBLE] be s2n minus 1 minus Cp times 2n S2.

IAN BALL: Yeah. I think 2n minus 1 because we started at 0. Well, let's go over it. Yeah?

STUDENT: That would be-- the 2n minus 2 is the 2n minus [INAUDIBLE].

IAN BALL: Yeah. So it'll be-- so I'll just write it, and then I'll explain. But yeah, it's easy. And I might get confused with some of these fence posts issues. But I think it's going to be Cp times 2n. Why is this?

Well, we're in period n minus 1. If I accept the offer, then I'm going to get a settlement of S2n and minus 1. But I've also had to pay my legal fees. I'm the plaintiff, so my legal fees are Cp per period. And we negotiated in period 0 and period 1 and period 2 and period 3, all the way up to period 2n minus 1. You might say that's 2n minus 1 periods. But remember, we started at 0. So we actually have 2n periods of negotiation costs.

And then what if I reject the offer? What's going to happen? Yeah?

STUDENT: J minus p hat.

IAN BALL: Right. And we can't forget-- I've already-- I mean, it's a sunk cost, but then we have this as well. So it's going to be J minus Cp hat minus cp times 2n. What does this mean?

Well, if I reject the offer, we're going to go to court. I get a judgment of J, as the plaintiff. I have to additionally pay all the court costs. And then I have all these costs from before.

But I think what people are kind of jumping ahead is these don't really matter. So we can circle these. At this point, these are what are called sunk costs. I've already paid all of these costs.

So in my decision today of whether to accept or reject, these aren't actually going to have any bearing on my choice because I pay them either way. I've already paid them. I've already paid the lawyers. The only choice I really have is, am I going to pay this additional trial fee, Cp hat? Great.

So now if we compare these two payouts, when should the plaintiff accept the offer? What kind of offers should the plaintiff accept? Well, intuitively, they're the plaintiff. They want more money. So we always think the plaintiff should accept-- let's just write this down here.

It's always going to take this form that they're going to accept, if and only if, or maybe I'll say use arrows-- if and only if S2n minus 1 is greater than or equal to some threshold. Let me denote this threshold as S2n minus 1 star.

And this is just-- we know this is generally how things are. If you're the plaintiff, this is how much you get paid. Getting offered more money is always better. So it makes sense that you're going to accept if the offer is high enough, and you're going to reject if the offer is too low. The question is, what is this critical threshold? What's the amount that you have to offer me as the plaintiff for me to accept your offer? Yeah?

STUDENT: J minus Cp hat.

IAN BALL: Right. J minus Cp hat. And you might say, well, we have to be a little careful if the offer is exactly equal to this. Maybe I'll accept. Maybe I won't. We're going to make the convention that when indifferent, I always accept. And this is what we have to do to make sure that equilibrium exists.

But notice what's interesting here. What is the plaintiff willing to accept? They're willing to accept some offers that are strictly less than the amount that they would get in court. So what does this reflect? This reflects the cost of going to court.

If you say, look, I'm not going to pay you as much as you'll get in court. I'll pay you a bit less. But it's still worth it for you because if you accept my lowball offer, you get to avoid the cost of going to court. So we can see the higher that Cp hat is, the weaker the bargaining position of the plaintiff, the lower the offer that they're going to be willing to accept because they really don't want to go to court and pay this enormous fee.

So now we go back one step earlier. We're still in period 2n minus 1. But now let's think about the defendant's choice. Now the question is, how much should they offer, anticipating that the plaintiff will accept their offer exactly under these conditions?

Well, intuitively, they want to make the smallest offer that will be accepted. They want to pay the least amount that they have to. So it's going to be optimal for them to choose S2n minus 1 equals S2n minus 1 star.

They could offer to pay strictly more than this. That will still be accepted, but that's worse for them because they don't want to pay out very much. I think it's a little tricky, though, because they might also want their offer to get rejected. So we still have to be a little careful. Let's see. What if the offer gets rejected?

We're claiming that they want to make an offer that will be accepted, but we actually have to see what happens if it gets rejected. It turns out, if it gets rejected-- and let's just work this through-- this will be accepted, and then they will get minus S2n minus 1 star minus cD 2n.

And what if they were to make some different offer that were rejected? What would their payout-- what would their payoff be? Yeah?

STUDENT: Minus J [INAUDIBLE] 2n minus C hat.

IAN BALL: Exactly. So if they make an offer that's rejected, then they're going to go to court. And that means they're going to lose J. They're going to lose CD hat, and they're going to lose CD 2n

But this is clearly worse. Because if they go to court, they have to pay more. They pay J rather than J minus Cp hat, and they have to pay the cost of court. So this is kind of extra bad. So they pay more, and they pay the court costs.

And we'll see throughout that it's never really going to be optimal for people to make offers that are going to be rejected. It's always just going to make things worse for them. So indeed, this is going to be the best choice for CD.

Now let's try to go one step back and see what's going to happen in period 2n minus 2. So maybe I'll come back over here. And then hopefully, we'll be able to see a pattern.

So now let's go to period 2n minus 2. And our conjecture is that the-- I have to do it right here-- the defendant is going to accept an offer of S2n minus 2, if and only if, S2n minus 2. Well, let's see. I'm the defendant now, so I'm willing to accept low offers. The plaintiff was willing to accept high offers. The defendant is willing to accept low offers.

So there's going to be some threshold, S2n minus 2 star, such that the defendant will accept if they're offered to pay a small enough amount. Is there an issue over here? Yeah?

STUDENT: Yeah, I have a question.

IAN BALL: Yes,

STUDENT: Basically, the comparison between [INAUDIBLE] giving an offer and being rejected. So getting an offer, they have to pay S2m minus 1 star, which is J minus Cp versus-- oh, J plus.

IAN BALL: Yeah. So we have to be careful with the minuses. So this is going to be minus J plus Cp hat.

STUDENT: So does it depend on how Cp is?

IAN BALL: It doesn't. And that's why-- and this is what I meant about it's kind of doubly bad for them. Because if their offer is rejected, they're going to have to pay the full settlement, J, in court. So going the defendant will have to pay more in court. And on top of that, they'll have to pay their legal fees.

So the difference between these two is actually the sum of Cp hat and Cd hat. It's not about the relationship between them. Yeah. Good point. So just to be clear what the argument is. The argument is, I could make some offer that's rejected. If it's rejected, the amount that I offer doesn't matter. All that matters, it's rejected. And we've shown that that's worse than making this offer.

The alternative is I can make an offer that will be accepted. But to make an offer that will be accepted, I'd have to offer even more than this. And that is certainly worse for me. So the best thing I can do is offer the smallest amount that will be accepted.

So now what our guess is that the structure in period 2n minus 2 is going to be something like this. The defendant is going to have some threshold where they're going to say, look, I'll accept the offer if and only if it's small enough, if and only if it's below some threshold.

And then let's think about the plaintiff's problem. We could go through the algebra to check. But you can see that the plaintiff wants the offer to be, the settlement to be as large as possible. But if it's too large, it won't be accepted. So they want to make the largest demand, the largest settlement that will be accepted. And that's precisely S2n minus 2 star.

So all we have to do is find this number. And that's how these games reduce. At each stage, one side is going to accept if and only if the offer is above or below some threshold. And the name of the game is finding these thresholds.

So let's see. Well, let's try to understand what this threshold needs to be, by let's-- what's going up here. So the threshold needs to be such that they're indifferent between accepting this offer and rejecting it. That has to be true at a threshold. Because if they weren't indifferent, there's no way that could be the threshold.

So if they accept this offer, what are they going to get? They're the defendant, so they have to pay it. So they're going to get minus S2n minus 2 star minus all the legal fees that they have. So that's minus CD times 2n minus 1.

So now we're one period earlier. So if we accept a settlement here, we don't have to pay legal fees in the last period. And this is how much we're going to have to pay. So we have a negative sign in front.

What happens if we reject this offer? Well, this is where backward induction comes into play. Because in general, I don't know what happens, what my payoff is from rejecting this offer, because it depends on what's then offered to me in the next period. But because we've already used backward induction in the next period, we know what's going to happen in the next period.

In the next period, the plaintiff is going to offer exactly this. And then it's going to be optimal for me, as the defendant, to accept it. So if I wait till the next period, I'm going to have to pay negative S2n minus 1 star minus-- what are my legal fees if I accept this in the next period? Yeah?

STUDENT: [INAUDIBLE] CD times 2n.

IAN BALL: Times 2n, right? Because if I reject today and wait till tomorrow, I have one more period of negotiation. So that means it's going to be CD times 2n.

And now we see the pattern. What are we going to find? We're going to find that S2n minus 2 star equals S2n minus 1 star plus CD. That's right. Yes, that's right, which we can do our math over here. And this is J minus Cp hat plus CD.

Why is it a plus? We could try to go through the algebra. But can we just reason through this, why we get a plus CD here? Why am I, as the defendant, willing to accept a settlement that's strictly higher than the settlement I would get tomorrow? Yeah?

STUDENT: Because you're not going to tomorrow. So you're not going to have the CD from the next period. So you're adding it back.

IAN BALL: Exactly right. So what am I saying? I'm saying, look, I can wait till tomorrow, and then I have to pay you this much. That's going to be the outcome of negotiation tomorrow. But if I pay you an amount today, I save one unit of negotiation costs. So I'm actually willing to pay more today, I'm willing to get a worse offer today because I'm saving the legal costs of going tomorrow.

And you can see how these enter differently. The plaintiff was also willing to accept a worse offer to save on legal costs. But for the plaintiff, a worse offer is a smaller settlement because they get paid the settlement. For the defendant, they're willing to accept a worse offer to save on legal fees. But for them, a worse offer means paying more.

So if I'm the defendant, I'm basically saying, look, I'll pay you off to save me the money of going to court. And the plaintiff is saying, you better pay me a lot. Sorry. I'll accept not paying you very much in order to avoid going to court. I'll accept you not paying me very much to avoid going to court.

So now I think we can keep doing this all day. But I think we're going to see a pattern emerge. So let's see if we can guess what the pattern is going to be for these thresholds. So it's always easy to get off by one. Let's see. I've done it right.

Let's look at the-- we started in an odd period. So we have the odd periods and the even periods. In the odd periods, the defendant is going to propose S. Let's say the odd period is 2n minus 2k minus 1. And in the even period, it's 2n minus 2k. And here, we're going to have k greater than or equal to 1.

So here, let's just write down the formula and then check that we have it right. I think the defendant is going to propose J minus. Cp hat plus k minus 1 of k CD minus CP. And here, I think the plaintiff is going to propose one step earlier. Let's check if I did this right.

So let's first try k equals 1. If k equals 1, we're in period 2n minus 1. So we better get J minus CP hat, And that's exactly what we get. If k equals 1 down here, we're in 2n minus 2, and we better get J minus CP hat plus CD. So think that looks correct.

And then let's just check what's happening here-- accepts if and only if, and here the defendant accepts if and only if. Well, in these periods they're going to accept if and only if. Well, the plaintiff is the one who receives the money, is only going to accept it if it's large enough, if and only if the settlement is greater than or equal to this number here. And the defendant is going to accept if and only if f is less than or equal to this number here.

So the defendant is proposing the smallest settlement amount that will be accepted because they like low settlement amounts. And here, the plaintiff is proposing the largest settlement amount that will be accepted because they like to receive high settlements.

So this is the strategies. We specified all the strategies at every contingency. But let's now compute what actually happens. The strategies say, in every contingency what do people do? But what actually happens in this big game? We made it really complicated. But what happens is actually going to be pretty simple.

You want to see what is actually going to happen? So when are they going to settle? Are they going to settle? Do we go to trial? What happens? We have the strategy here. We just have to figure out what's actually going to be done.

Well, let's start in the first period, and let's go through it. In period 0, what is going to happen? Well, period 0 means k equals n. So the plaintiff is going to propose this settlement amount, J minus Cp hat plus n minus 1 times this plus Cd. And then the defendant's going to accept it.

So that's it. So the outcome is actually very simple. The outcome is we settle in period 0 for this amount down here. Yeah, this looks right. I'm just making sure I'm not off by one index.

So can we try to interpret this amount? How does this amount compare to J? Well, we see that the settlement is going to be smaller if Cp hat is really big. So this settlement is going to tend to be smaller if Cp hat is big. And it's going to tend to be bigger if Cd minus Cp is large.

So do we have intuition for this? Well, let's see. If Cd minus Cp is large, that means the defendant has higher negotiation costs. And that means the defendant is at a strategic disadvantage. The fact that they have high negotiation costs means the other side can drive a hard bargain. And because of that, the defendant is going to have to end up paying a larger settlement.

Notice that it's not only big if Cd minus Cp is big, but it also depends on how far away the trial is. So if the trial is really far away, then the defendant is really in trouble because the defendant's higher legal fees per period could really add up over the many, many periods that going to trial.

On the other hand, if the plaintiffs cost of going to trial is really, really high, then the settlement is going to be smaller because this is going to advantage the defendant. They're not going to have to pay out as much if they can drive a hard bargain to the plaintiff in the last period.

But here's one kind of puzzle. What happened to Cd hat? It doesn't appear at all. So why does it not depend at all on Cd hat? Yeah?

STUDENT: I'd say that the plaintiff is always going to accept on the last day, so the cost of trial is never realized.

IAN BALL: That's right. But we also never get to the fifth period. But the costs that we would incur if we get to the fifth period also show up. So I think that's part of the way there. Yeah?

STUDENT: Because the defendant is the last proposer. So they can base it on what the plaintiff will accept.

IAN BALL: Exactly right. So if we actually change the timing here-- and this is one issue with these bargaining games that sometimes the details matter a lot. If we had an odd number of periods before trial, this would kind of flip. And we would care about Cd hat, not Cp hat.

And the reason is that the defendant is the one able to exploit the plaintiff's cost of trial because the defendant gets to make the last offer before trial. So because the defendant is the one who makes the last offer before trial, the defendant can say, you better give me a really good deal. Otherwise, we're going to have to pay this high cost to go to trial.

It's interesting, though. I mean, the defendant also pays a high cost to go to trial. But they're able to exploit the rationality of the plaintiff and make them an offer that they're going to have to accept, and therefore, take advantage of their bargaining power. Great.

So I guess a few takeaways-- this model suggests that we shouldn't have drawn-out negotiations, that the settlement should end right away. But notice that the potential length of the negotiation affects the outcome. So even though we settle in the very first period, the potential to have a very long, drawn-out negotiation drives the amount that we settle for in this first period.

One final thing. If on an exam you're asked about this, the strategy is not we settle in the first period. The strategy is this really complicated thing. The outcome is that we settle in the first period.

Now, in reality, we often see long drawn-out negotiations. So I think it's often very helpful in economic models to say, what's missing? Why is it that in reality-- let's say Trump and the Democrats, I think they're going to be negotiating for a pretty long time would be my guess. That's completely consistent with the prediction of this model. So what do you think is missing in this model that explains why everything ends in the first period? Yeah?

STUDENT: Uncertainty about what J actually is.

IAN BALL: Exactly. An uncertainty about the preferences of the other side, what they're willing to accept. Here, we all know exactly what the other side is going to accept, and we can push them just to the brink of acceptance at every round. And therefore, there's no reason for things to evolve over time.

The reason that we often have delay and impasses in bargaining is that we don't know what the other side is willing to accept, and we might want to make them a really bad deal early on on the off chance that they end up not being who we thought they were. And we'll see later in the course models of bargaining with incomplete information or uncertainty about the opponent's preferences. And that will exactly lead to delay in bargaining.

Imagine you were negotiating with someone, and you didn't know if they were a tough negotiator or a soft negotiator. You might want to make a really ridiculous offer in the first period that only the soft negotiator would accept, but that you know the hard negotiator would not accept. And then move to the next period and change your strategy. So that's the basic idea of what goes on when you have incomplete information about the other players.

OK, let's now move to the second bargaining situation today, which is going to be haggling over the price. So now we're going to imagine that-- we'll call this price haggling. Let's suppose that we have a seller and a buyer. And the seller just wants to get paid. They're negotiating over the price of some item.

So if they trade at price p, then the seller, their utility is going to be the price p that they get. The buyer's utility is going to be v minus p. So that means p is the price, and v is the buyer's value for the good.

So notice here, we're assuming that it doesn't cost anything for the seller to produce the good. That's a simplifying assumption. It's already been produced. Maybe they can't resell it to anyone else. They're just stuck with it.

So the seller values the good at 0. They're thinking of selling it to the buyer who values it at v, and they're trying to haggle over what the price should be. So first of all, let's just make sure we understand what's a reasonable range of what the price can be? So I think there's some price that we can rule out.

How low could the price be? What's the lowest it could possibly be? Well, actually, it could be pretty low. I think it could go all the way down to 0, in principle. It's certainly not going to be less than 0, though. Because if it was less than 0, well, the seller would never accept being paid a negative price.

What about the highest price that we would expect to see here? Yeah?

STUDENT: v.

IAN BALL: v. It certainly can't be more than v because the buyer would never accept a price above v. So we can, at first pass, just by basic individual rationality concerns, see that the price is going to be somewhere between 0 and v, and both players are going to get some non-negative utility. But the question is, where is it between 0 and v?

The seller would love it to be very close to v, and basically leave the buyer with nothing and charge them their full valuation for the good. The buyer would love to pay a price that's really, really low, so that they can get all the value from the good. And we're going to try to make some predictions about where within this interval the price can be.

Notice this is kind of a common situation that we generally only expect that goods are going to be exchanged when the seller values the good more than the buyer. And we generally expect the price to be somewhere between-- sorry-- when the seller values the good less than the buyer. And we expect the price to be somewhere between what the seller values it at and what the buyer values it at. And here, what the seller values it at is 0. You can think of this as the seller's value, and this is the buyer's value.

So the buyer does value it more. That's why it makes sense for them to buy it. But where the price lies between these two valuations is the name of the game. That's what we're trying to figure out.

Now, what we're also going to imagine is that there's a cost of delay. But we're going to model the cost of delay a bit differently. Here, we model the cost of delay as an additive cost. Each period we had to pay a lawyer some daily fee, and that had a cost. Here, we're going to say the cost of delay, we're going to use what's called discounting.

So the idea is that if we trade in period t, then the utilities we actually get would be delta p for the seller, and delta v minus p for the buyer to the t, where delta is between 0 and 1. So here-- you may have seen this in other economics courses. Delta is called our discount factor. And what this says is each period we wait our utilities get multiplied by this number delta, which is less than 1. And therefore, they get smaller.

So if t is really, really big, then the seller doesn't get the full utility from the price. They get that times delta to the t because each period they have to wait beyond the zeroth period they incur this multiplicative discount of delta. And similarly, for the buyer. They incur this multiplicative discount of delta.

So let's understand if delta is very close to 1, then that would indicate that the players are quite patient. Because as time passes, delta to the t doesn't get smaller that quickly if delta is very close to 1. If delta is very close to 0, the players are extremely impatient. And as time passes, their utility is going to drop a lot. And therefore, the cost of delay becomes very high.

So maybe I'll-- you can say this in a few ways. You might say low delta is another way of saying the players are impatient. Or another way of saying that is the cost of delay is high. And then conversely, if we have high delta, the players are very patient, and the cost of delay is not very high.

Sometimes we think of delta as being related to the interest rate. So in a financial context, the cost of delay with money is the interest rate that you could have gotten by putting that money in a bank account or getting some return in equities. And now we're going to assume that time proceeds just as before.

So I think this is kind of exactly what you see if you go a market, especially in developing countries and there aren't posted prices. You, as the buyer, go up and you say, I'm willing to pay you this. And they say, oh, no. You need to pay me this much. And then you counteroffer, and you go back and forth.

And even in developed countries, this certainly happens if we're negotiating on some big deal. It doesn't happen at Starbucks. You're probably not going to be able to negotiate over the price of your coffee. But if you're getting hired at a firm, or if you're making a merger, people do negotiate all the time on prices.

So 2n minus 2, 2n minus 1, and then 2n. And the order I think we're going to assume is the seller makes the first offer. So here, the seller offers a price, P0. So instead of offering a settlement amount, they're offering a price. The buyer can either accept or reject.

If it's accepted, we're done. If it's rejected, now the buyer gets to make a counteroffer of P1, and the seller can accept or reject. And we're going to go on like this, all the way up to seller making an offer P2n minus 2. Buyer making an offer P2n minus 1.

And then we're going to assume at period 2n, the world ends. Or just, that's it. We both get 0. So if we wait this long, we get payoffs of 0, 0. This seems like a strange assumption.

The real reason is that we don't quite have the tools to analyze what happens if it could go on indefinitely. But it turns out that if you take n to be really, really big, the predictions that we'll get will be very similar to the predictions if there's an indefinite end date.

So as I tell the story of the tourists get off the tour bus at the market, and the tour bus is leaving in 2n periods, and this is the last chance for them to make a trade. But are they really-- or people say the ice cream cone melts. But yeah, I don't know how compelling this exact end date is, and we'll talk about relaxing that assumption later on.

So now let's try to start analyzing this. So let's start, as usual, in period 2n minus 1. And now what is the seller going to accept?

So we're in period 2n minus 1. The world is going to end tomorrow. The buyer makes an offer. What prices will the seller be willing to accept? So they'll accept if and only if the offered price, P2n minus 1, is greater than or equal to what? Any thoughts?

So you're the seller. The buyer says, I'm going to pay you this amount for it. You say, well, either I can accept this, or I can reject it and get a payoff of 0. So what am I willing to accept? Yeah?

STUDENT: Won't it just be 0?

IAN BALL: Yeah, anything that's above 0. Exactly 0, we can argue about. But we'll say if and only if it's greater than or equal to 0. So if you offer me any non-negative price, my choice is I accept it, or I wait till tomorrow and I get 0. So any positive amount you give me, I'm willing to accept because it's better than getting 0 tomorrow.

And given that, what is the buyer going to offer? So again, we're working backwards. What will the buyer offer? How much will they offer to pay? It's a good position for the buyer, right? They're going to offer to pay 0.

If I know that the seller will accept any price, I'll offer the smallest price they'll accept. What's the story here? I come at the end of the day, think about-- I mean, we see this all the time. It's a store. They have one croissant left. They're about to close. Or maybe they have tons of croissants left.

I know if they don't sell it to me, they're not going to sell it to anyone. It'll go stale tomorrow. I can make a pretty good-- I'm in a pretty strong bargaining position. I can say, just give it to me for a price of 0. And we see bakeries have discounts at the end of the day for exactly this reason.

But now let's go to period 2n minus 2. Now the buyer is making an offer. Sorry. The seller is making an offer. So the buyer is choosing whether to accept the offer or not. And what offers will the buyer accept? So the buyer will accept if the price, pt, offered by the seller is small enough. Yeah, question or-- yeah?

STUDENT: Is the buyer's value of the good known to the seller?

IAN BALL: It is. And that's a crucial assumption. So it is known. And in reality, you might think it wouldn't be known. And towards the end of the course, we'll talk about incomplete information bargaining. But it comes really messy. But yes, v is known. That's a good point. And delta is also known.

So here, we're going to have, say, Pt star. The buyer will accept if the price is low enough. And let's try to compute what Pt star has to be. So as usual, Pt star is going to be defined by an indifference condition. The highest price they'll accept is the price where if they accept it, they do just as well as if they were to reject it and wait until tomorrow.

So let's look at what happens if they were to accept a price of Pt star, then they're going to get delta to the 2n minus 2 of v minus Pt star. So that's if they accept.

What if they reject the offer? If they reject the offer, then tomorrow they get to propose, and they're in a great shape. We already said tomorrow they're going to offer a price of 0, and it's going to be accepted.

So that means tomorrow they're going to get delta to the 2n minus 1 times v. Really, it's minus v PT plus 1 star. Or I guess I should say, here, t equals 2n minus 2. But I'm just writing t because it's a bit cleaner.

And once again, we can see that we can simplify these discount factors a lot. Because we can factor out delta to the 2n minus 1 from both terms. And this makes sense because the question is, do I take it today, or do I wait till tomorrow?

There's really only one discount factor that matters. It's the discounting between today and tomorrow. So we can equivalently-- we can cancel things out, and we can say-- and just put a delta here. So we divide by delta to the 2n minus 2. So what we get is v. And maybe I'll write the general formula-- v minus Pt star equals delta.

In period t, in an even period where I'm the buyer, I can either get this utility today, or I can wait till tomorrow and get this utility where the price I get tomorrow is going to depend on what we've already computed for tomorrow using backward induction. And if we try to solve for this, we'll get Pt.

Now, let's ask, how does this Pt star, this threshold price, compare to Pt plus 1? Is it bigger or smaller than Pt plus 1? Well, it's bigger because Pt is a weighted average of Pt plus 1 and the number v. But we know v is always bigger than or equal to Pt plus 1. So that means this is strictly greater than Pt plus 1. And let's try to think through the intuition for this.

Again, it's about discounting. I'm willing to accept a price that's higher today than I would get tomorrow for the classic reason that if I wait till tomorrow, I'd have to pay a cost of delay. So I'm willing to accept a worse deal today than I would get tomorrow in order to avoid the cost of delay. For me, a worse deal is a higher price because I'm the buyer.

So I started out with a case of 2n minus 2, but you can see that this formula actually holds for any even period. For any even period t, the buyer is going to accept if the price is below a threshold. The seller is then going to offer exactly that threshold. And the threshold is defined by this formula.

Now, let's try to look at odd periods t. If we're in an odd period t, now it is the seller who's deciding whether to accept or reject. The seller is going to accept if Pt is greater than or equal to Pt star.

And now can we get a formula for what Pt star is going to be? It has to say that the seller is indifferent between selling at a price of Pt star today or waiting until tomorrow. If they sell at Pt star today, they get Pt star in discounted terms. What happens if they wait till tomorrow?

Well, they get delta times Pt plus 1 star. If I wait till tomorrow, we know from backward induction that the good is going to be exchanged at a price of Pt plus 1 star. But I have to discount that by delta, and I'm going to be exactly indifferent if this holds.

So if you want to be kind of really mathy about it, we have a solution, but it's a bit of a mess. So what we know-- so how do we compute our solution? We just work backwards. We know-- we argued that the price that's going to be determined in period 2n minus 1 is exactly 0.

And then we have a sequence of recursions that allow us to compute what Pt must be if we know Pt plus 1, what Pt is if we know Pt plus 1, and we have two different formulas depending on whether if t is even or odd.

So technically, we could start at the end and work our way backwards. Then we're going to have some really messy formulas. It's going to get ugly. So I'm not going to go through all the algebra, but I just want to show you what we're going to find, so what's the limit of this game.

So we can compute the full strategy. The outcome is going to be-- well, in period 0, the seller offers P0 star, and the buyer accepts, where P0 star is determined by starting at the end and working all the way backwards. So once again, we see no delay. And this already came up earlier.

What's the reason that we have no delay? Why do we not see this protracted negotiation over the price that's going to be paid? What's kind of missing from the model? I think someone already brought it up. Yeah?

STUDENT: They have perfect information so they can see what happens in there.

IAN BALL: Exactly. In reality, when you're negotiating with a price for someone, you don't know their value, v. You don't know exactly how much they're willing to pay. We've made this really stylized assumption that everyone knows the buyer's willingness to pay-- exactly the most the buyer is willing to pay for this.

That would make-- say you're negotiating over a house. It would be really easy if you understood exactly how much this family was willing to pay to buy this house. In reality, we don't know that, and that's why we see delay. But in this benchmark model, with complete information, we see no delay.

I think this is an example of where an economic model can be useful, even if it gives us results that are counterintuitive. Because what it can show us is it can tell us-- demonstrate the importance of certain features. What this tells us is that the driver of delay in bargaining is really incomplete information about people's preferences.

How do we see this? Well, we see if there were no incomplete information, we would have no delay. And this shows us that something must be missing from the model. And I guess we suspect it's incomplete information. And we can see that the later. We can model that later.

I want to describe, though, briefly, what's the limiting case if n is very large? And I explained this is what we're going to converge to if we looked at the infinite horizon model. And it turns out that P0 star is going to be 1 over 1 plus delta times v.

Let's try to interpret this. And this is the case-- I should say in our model, the seller moves first. What if we had said had the buyer proposed first? We can also compute what P0 star is. Any guesses if the buyer--

So we're still looking at the limit where we could, in principle, negotiate for a very long time. This is the price that the seller is going to offer if they propose first. What if the buyer got to propose first? Any guesses on whether this price would be higher or lower?

Again, we can do the algebra, but I think it's a lot better to think through these things intuitively. Would you rather go first or not in this negotiation? Any intuition? It's not so easy?

It turns out going first is generally an advantage. And the reason is that if I go first, I can take advantage of the fact that, yes, we both experience a cost of delay. But, in particular, the other party experiences a cost of delay. And therefore, they must be willing to accept offers that aren't that good for them.

So if I go first, I can say, look, I'm going to give you a pretty bad offer, and I know you're going to have to accept it. Because if you don't accept it, you're going to have to pay the cost of delay. And this gives kind of a first mover advantage. And it turns out what we're going to get here is P0 star is going to be delta over 1 plus delta times v.

So if the buyer goes first, the price the buyer is going to get is going to be strictly smaller by a factor of delta relative to what the seller would get if the seller went first. And here, we can see the role of discounting.

What happens if delta goes to 1? If delta goes to 1, the price the seller gets if they go first converges to v over 2. And the price the buyer gets if they go first also converges to v over 2. So we see as delta goes to 1, and the players become very impatient, the first mover advantage disappears.

Can someone see-- if they can provide intuition, why does the first mover advantage disappear if delta goes to 1? Yeah?

STUDENT: Just to make sure, when delta goes 1, they both are very patient [INAUDIBLE].

IAN BALL: Exactly right.

STUDENT: [INAUDIBLE] if they don't care about waiting [INAUDIBLE].

IAN BALL: Right. If we get very, very patient, then it shouldn't really matter who goes first. Because tomorrow the other person gets to go first. And tomorrow isn't that much different than today if delta gets very close to 1.

Remember, another interpretation of delta going to 1, one interpretation is the players are literally more patient. Another interpretation is they can make offers in very rapid succession. So it could be that the reason period 1 is not very discounted relative to period 0 is that the offer made in period 1 happens right after the offer in period 0. They happen in really quick succession.

And intuitively, if we're making alternating offers every millisecond, it shouldn't really matter who goes first. And that's exactly what we see. The first mover advantage disappears.

Let me just to conclude, try to derive, give a heuristic derivation of these limiting formulas, and then we'll conclude. So let's-- so how can we try to solve for, let's say, the first formula?

So let's say the seller proposes. Then what we expect is if the seller proposes, and we're taking the limit as the number of periods get very large, then the difference between period 0 and period 2, we think that it should get very, very close.

Because in period 2, we have a billion periods to go. In period 0, we have a billion plus 2 periods to go. Since it's a billion, they should be pretty similar. This is a heuristic derivation. And now let's plug this in and see what we can get.

So we get P0 star from this recursion is equal to 1 minus delta times v plus delta times P1 star. But then using this formula, we know that P1 star equals delta times P2 star. I'm just applying this formula with t equals 0, and then I'm applying this formula with t equals 1. So now let's plug this into here, and we get P0 star equals 1 minus delta v plus delta squared P2 star.

Now, in general, P0 star and P2 star are going to be different. But in the limit as n goes to infinity, P0 and P2 star are going to look really, really similar. And the reason is that, again, if there's many, many periods ahead of us, a billion periods and a billion plus 2 periods, it should be basically the same. So let's squint and make this P0 star.

Now let's try to solve this. We bring this to the other side. So we get 1 minus delta squared P0 star equals 1 minus delta v. And now there's one kind of trick that I think is kind of good to remember on problem sets. It's not so clear if we solve for this, we get P0 star equals 1 minus delta over 1 minus delta squared times v. Uh-oh. That's not the formula we got, but it simplifies. Yeah?

STUDENT: Factor.

IAN BALL: We can factor, right? 1 minus delta squared is just 1 minus delta times 1 plus delta. The 1 minus deltas cancel, and we exactly get 1 over 1 plus delta times v. So again, I don't want to make the problem set for the exams all about doing algebra, but this is one trick that I think is good to know.

Let me stop there-- a little bit early, but I think we're good. And we'll continue on Tuesday with subgame perfect Nash equilibrium. I'll see you then.