[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: So if you recall last class, we discussed finitely repeated games. Today, we're going to discuss infinitely repeated games, and we're going to see that the results are actually quite different.

So what's the motivation for studying infinitely repeated games? We can, of course, debate about whether the future is really infinite, whether it's finite. We can have these philosophical debates, but I think what we're hoping to capture with these models is simply the idea that there's not a natural last period. So there's no natural last period.

So when firms are setting prices, there's never really a time where they think there's no opportunity to set prices in the future. When governments are negotiating, there's no natural point where they think the world ends and we're not going to have further negotiation.

And what we saw is that when we modeled finitely repeated games, a lot of the conclusions were driven by these last period effects. By the fact that in a finite repeated game, there's some period of the game that everyone knows is the last period of the game, and the game cannot continue beyond that. We want to shut that down today and study games where there's no natural last period.

There's going to be some key differences between the finite and the infinite case. And what we saw in the finite case is in a lot of games, like in the prisoner's dilemma, there was no scope for cooperation or for punishments and rewards. So maybe I'll say limited scope for rewards and punishments.

And this came from this last period effect. We often argued if we know what's going to happen in the last period, then we can figure out what's going to happen in the period before, and we can keep working backwards using this kind of backward induction reasoning. And it's all driven by the fact that in the last period, there's certainly no scope for rewards and punishments, and then that's extended earlier in time.

In the infinite case, we're going to see, there's actually a huge range of possibilities. So maybe we'll say many things can happen. We'll actually formalize this next class, but particularly if players are very patient. We're going to see that if the players are patient enough, the future looms very large, and there's a huge richness of possibilities for equilibria in these games.

So today, the plan is to-- we're not going to quite get to this anything can happen result until next class, but today, we're going to set up the framework of infinitely repeated games, and then we're going to discuss how we can achieve cooperation in the prisoner's dilemma. So we're going to focus on the prisoner's dilemma, and we're going to observe a crucial difference in the prisoner's dilemma, the unique subgame perfect Nash equilibrium was defection every period, we're going to see in the infinitely repeated version, there's scope for a lot of cooperation.

So let's set up the general framework. So just as in the finite repeated version of the game, we have to have some stage game that we're going to repeat. So as before, we have a stage game. And maybe we'll call this G. Remember, we had our stage game G, and this was just a simple strategic form game, but remember, we defined or labeled strategies in the stage game as actions rather than strategies so that we don't get confused.

So this was-- we had ui from A to R for i equals 1 through n. So we have n players. The player's payoffs depend on action profiles. So remember, A is A1 cross up to An. It's the set of action profiles the players can play in the stage game, and then we specify the utilities that the players get as a function of that.

Now the difference here is timing. Whereas before, we said that time was 0, 1, 2, up to some fixed finite horizon, now the horizon is infinite. So now we have periods t equals 0, 1, 2, dot-dot-dot. There's no point at which the game necessarily ends. And otherwise, everything's the same. So in particular, we still have perfect monitoring of past actions.

Now the first thing we have to figure out in a game like this is how we're going to compute payoffs because in the finitely repeated version of the game, we just could add up the payoffs and then take the average. But here, things go on forever, so it's not clear how we're going to be able to add up all the payoffs. And we're going to return to an idea that we had before, which is we're going to have discounting.

So we're going to use discounted payoffs. So let's see what we mean by that. Well, what's actually going to happen? What's an outcome of the game? In the finite version, an outcome of the game was a profile of actions in the first-- in the zeroth period, a profile of actions in the first period, all the way up to the last period. But now, an outcome of the game is an infinite path of actions.

So we might say, we want to-- be interested in something called an action path, which is something like a 0, a 1, a 2, dot-dot-dot. So this is like an infinite vector. And this is going to say, well, what did everyone do in period 0? Notice, there are no subscripts here. If I had a subscript i here, that would mean what player I did in period 0. But with no subscript, this is what everyone does in period 0. It's a profile of actions. So just to make sure we're on the same page, each of these lives in A.

So an action path says, what did everyone do in period 0? What did everyone do in period 1? What did everyone do in period 2? And so on. And what we want to understand is what payoffs the players get from an action path like this? You could think of the action path as playing the role of the terminal nodes in a finite horizon game.

Remember, we wrote out the extensive-form game, and we had this tree, and we looked at the terminal nodes. Here, the game goes on forever, so it's a bit trickier to think about what a terminal node is, but this is kind of the analog of a terminal node in the infinite horizon game.

And now we need to understand what payoff this gives. So the payoff from this is going to be-- maybe we'll write ui of a0, a1, dot-dot-dot, we need to say what this is. And again, we can't just add it all together because it's going to go on forever, so we're going to look at the discounted sum of the payoffs.

So player i will get ui of a0 in period 0 because this is the action profile that's played in period 0, and this is the stage game utility that player i gets. So you can think of this as player i's flow payoff in period 0. Then we want to add to that the flow payoff in period 1, but now we have to discount that.

So we're going to get delta times ui of a 1, So this is the flow payoff player I gets from the action profile played in period 1 multiplied by delta, that's our discount factor, which we've looked at before, to represent the fact that payoffs tomorrow matter less than payoffs today.

And then we keep doing this. Then we'll get delta squared u of a 2, and we'll continue on like this. And if you want to write this a bit more mathematically, we can use summation notation, so it's the sum. We're summing over all the periods, and our index is starting at time 0. So we're going to have a sum from t equals 0 to infinity.

And then what is the discounted flow payoff in period t? It's delta to the t times ui of a t. Where, notice, when t equals 0, this just reduces to 1. Delta to the 0 is equal to 1, and we just get this first payoff here.

So I think this raises a few questions about the discount factor delta. The discount factor delta is going to be a number between 0 and 1. The idea is that the future matters less than the present. If delta equaled 1, then we wouldn't be able to compute the summation because we'd be summing infinitely many things, and that might not converge.

How do we interpret this discount factor delta? I don't know, how do you think of this? How can we think of this discount factor? Why might you value future payoffs less than payoffs today? Do you feel that you have a discount factor? If I said, I'll give you $10 today or $10 in 10 years, which would you prefer? Yeah?

AUDIENCE: Well, $10 today due to inflation.

IAN BALL: Sorry, say it again?

AUDIENCE: $10 today due to inflation.

IAN BALL: Inflation, right? Exactly. So one interpretation discount factor, which is especially relevant if we have monetary rewards, is inflation and interest rates. So you could take that money today and you could invest it in, say, Treasury bills, and you could get a guaranteed return to the future. So one interpretation is sometimes the interest rate, which is tied to inflation.

And under this interpretation, we might say that delta is equal to 1 over 1 plus r. So if we have an interest rate r, and we're thinking about monetary payoffs, then the value of $1 today is-- the value of $1 tomorrow has the same value as 1 over 1 plus r dollars today.

Why is that? Well, if you give me 1 over 1 plus r dollars today, and I invest that in a checking account, and I get a return at rate r, then I'll multiply this by 1 plus r, and I'll get my dollar back tomorrow.

Of course, sometimes these utilities are not monetary payoffs. So if this is just how much I enjoy a certain item, we can't literally-- I can't take my utility to the bank and get an interest rate. So sometimes we have other interpretations. Any other reasons why, if I said-- what if I told you I'm going to give you a really big reward in a billion years? Yeah?

AUDIENCE: We'd be dead by then.

IAN BALL: Right. So it could be people die, or maybe another way of saying it is, things could end, the game could end. There could be some reason why you leave the game, things could happen. So one other interpretation of this discount factor is that there could be some fixed probability at which the game ends each period. So if we wanted to formalize that, we could say, what if the game ends with probability 1 minus delta each period?

Well, then we can say, what's the probability that we actually get to tomorrow? Well, there's probability 1 minus delta that the game ends today. So the probability that we actually get to tomorrow is going to be delta. But then what's the probability that we get to the day after that? Well, again, there's only probability delta that the game continues.

So, one interpretation of delta to the t is saying in order for us to get to period t, we have to transition from perioc 0 to period 1 which has probability delta. We have to transition from period 1 to period 2, which is probably delta, and so on. And delta to the t is simply representing the probability that we're still around or that we survive up until period t. And crucially, if the game ends with probability 1 minus delta, then it continues with probability delta, so that's a key observation.

But even if you were certain the game would continue, I think people are just impatient, and this is something that's been documented in the lab all the time. If I say I'll give you an ice cream cone today or tomorrow, we don't have to worry about inflation because I'm giving you an ice cream cone, I'm not giving you money to pay for the ice cream cone, you're pretty likely to be around tomorrow, but I think a lot of people would just say, I'd rather have ice cream cone today than tomorrow, this is just a behavioral fundamental preference.

So we might just say-- maybe I'll call this fundamental impatience. People are just impatient. This is the way people are. In reality, you can think of this discount factor delta as maybe reflecting some combination of these forces.

Depending on the application, it may be really about interest rates, it may be about the game ending. It may be some combination. Maybe you're worried that the game will end before tomorrow. And you're also just fundamentally impatient. We're not really going to take a stance, and the mathematical conclusions won't really depend on the interpretation, but we're just going to have some fixed discount factor delta that's reflecting any reason for discounting. All right.

And because of this discount factor delta, sometimes this game-- remember, before we referred to a game-- so just a reminder, remember, when we did finite repetition we talked about G of T. So we said in the finitely repeated game, we denote it by G of T where we said this means game G is played up to Horizon capital T.

In the infinite case, we might talk about G of infinity, or more commonly, G of delta. Because just telling me G of infinity doesn't give me enough information. That tells me the game is repeated with an infinite horizon, but if I don't know the discount factor, I haven't really fully specified the game.

So more commonly, we'll write G of delta to indicate the infinitely repeated game with discount factor delta. More generally, we could have G of T, comma, delta, but normally in the finite case, we don't do discounting. So that's something to keep in mind.

All right. So now let's think about what a strategy is. So as before, we just define a strategy to be a mapping from histories to actions. So what are the contingencies you face in the game? You're a player in the game. In every period, you can see what happened before, and based upon what's happened so far, you choose an action.

So a strategy for player i is going to need to say, let's look at all the periods, period 0 all the way down to period t I'll do period 1 as well. Well if it's period 0, nothing has happened so far. There's just a null history. So we're just going to specify a single action. We need to say what is si of the empty set?

So this is going to say, my strategy needs to tell me what I'm going to do in period 0. I know I'm in period 0, nothing has happened so far, and I just need to choose an action.

Then in period 1, I don't just specify one action, I need to specify a lot of actions because I need to say what I'm going to do in period 1 as a function of or contingent on how we all played in period 0. And that's what h1 represents. So h1 is going to say in period 1, what history of play-- what is the history of play?

So for all-- well, where does h1 live? Let's call H 1. This is the set of histories at period 1. But can we be more precise? What is this set going to be? If it's period 1, and I observed what happened last period, what is H1 going to be? What's a history in period 1? Yeah, in the front.

AUDIENCE: What happened in period 2.

IAN BALL: Right. And what's our notation for that? Do we have-- so how can we describe what happened in period 0? What do we need to say-- or just in words? What do we need to specify from period 0?

AUDIENCE: I guess what each player did.

IAN BALL: What each player did. So we have to specify what player 1 did, what player 2 did, all the way to player n. That means we need to specify a profile of actions from period 0, and we have notation for that, that's just A. So remember, our set A contains profiles of actions that specify what every player does in the stage game.

And then if we go down to period t, now things get really complicated, because in period t, we have to specify what player i does as a function of every possible history of play. Now the history is a lot more complicated. So we have for all histories h t and H t equal-- so now what is h t? Well, it's period t, so I know how people played in period 0. I know how they played in period 1, and I know how they played all the way up to period t minus 1.

So that means altogether, I know t action profiles. So this is going to be At. So let's just make sure we understand. So as an example, h t would look something like this. It's going to say this is how people played in period 0, this is how people played in period 1, and this is how people played in t minus 1. So at time t history, looks something like this. Because we specify action profiles from 0 all the way up to t minus 1, we have to specify t action profiles altogether, and that's why we have A to the t here. So this set A to the t is just the set of all vectors that look like this. Yes?

AUDIENCE: And so little a0, that includes an action for every player?

IAN BALL: Exactly right. So this is kind of tricky because this is like a nested vector. What is a0? This is a 0 1 an 0. So you could if you wanted, you could think of this kind of like a matrix. We have to specify what every player did in every period. So we have to specify a number for each player and for each time. And we're going to be really consistent about this. The superscripts are always going to be about time and the subscripts are always going to be about the players.

So we have subscript i tells me which player I am I talking about, and that i can vary from 1 to n because those are the players. A superscript t can vary from 0 to infinity because that tells me what period I'm in.

Now I'd like this notation. It's consistent, but I guess one downside of this notation is it can be unclear whether a number like this is a superscript, just a label, or whether it actually means multiplying things together. So let me just give you a warning about this. When I write-- and if you're confused at any point, let me know. When I write delta to the t, delta is a number.

So here, delta to the t is the number delta multiplied by itself t times. So this is multiplication. When I write a 0 or a 1, I can't multiply a by itself because is not a number. This is just a symbol, it's just notation. And I guess I'm actually being a little, I guess, unclear here, a little naughty because this is just a symbol, this is just a symbol, but this actually is the product of sets.

So maybe I'll write-- to avoid that confusion, let me write this as A times itself t times. But if you're ever unsure about what notation means, let me know because I think sometimes the biggest barrier in studying repeated games is just getting notation right. But once you get the notation, I think everything should be clear.

OK. So formally, what is a strategy? Well, a strategy has to say what I do at every single history. So we could try to think of a strategy writing out these spaces and talking about, OK, we have the null history, we have some history here, we have some history here, and so on. The problem is this list is going to go on forever. So this is going to be pretty complicated to write down.

So we're just going to denote a strategy as a function. What is it? It's a function si from H to Ai. This says, at every history H, I'm going to specify what I do. But I haven't defined what the set H is, so let's just think carefully about what H is where H is just going to collect all the histories together. I'm just going to take all the histories and put them together. So I'm going to collect the period 0 history, well, that's just the empty set, so I need to put that in there. Then the period 1 histories. H 1 union H 2 union dot-dot-dot. Where's the eraser?

So here, the union-- this is just set theory notation, but it doesn't matter. All I'm saying is, I'm going to collect together the zero history, all the histories in period 1, all the histories in period 2, all the histories in period 3, and so on, and put them all together in a set. And then my strategy is going to tell me what action I take as a function of every possible history that I'm at. All right, any questions on this? Yeah?

AUDIENCE: So OK, just to make sure--

IAN BALL: Yeah, absolutely.

AUDIENCE: If you had a set of different actions, you would do for each possible history at each stage--

IAN BALL: Yeah, not-- so it will specify a single one. Ai says what I specify lives in Ai. So Ai is a set. Ai is the set of actions I could take. But a single strategy si, a single function will say at every single history, what action will I do?

AUDIENCE: Like, every single possible history for each period?

IAN BALL: At every single possible history for each period. So if I want to break it down, I could think of si as just encoding a different sub strategy for every history. It has kind of a period 0 part that tells me what I do in period 0 as a function of-- well, in period 0, there's only one history, so it's kind of easy. And period 1, what I do at every period 1 history. Then in period 2, what I do at every period 2 history, and on and on. So you think of it as a list of functions.

We're not really going to deal with this formal notation too much, but I just want us to know what the objects are and what's going on. Yeah, great question. Any other clarifications or questions? OK.

So now, we want to apply the one-shot deviation principle to understand subgame perfect Nash equilibria of this game. So now our goal is to study or analyze subgame perfect Nash equilibrium of these infinitely repeated games.

Now remember, step 1, if we want to analyze subgame perfect Nash equilibria, is always understand the subgames. So what we want to do is understand, what are all the subgames in this infinitely repeated game? And the key observation is that each subgame, each history h starts a subgame. But now what's tricky is the finitely repeated version, what the subgame looked like depended on how far along in the game we were.

If our subgame started right before the end of the game, we didn't have many periods left. But in the infinitely repeated game, whatever period we're in, the future actually looks the same. The future is still infinite. So let's try to understand what a subgame starting at history h looks like.

So just to make it maybe concrete, let's take t equals 2 just to make things really simple. So let's look at a history, and consider a history h 2, which is going to specify a 0, a 1. So what I'm going to say is let's look at the subgame where we're in period 2, we know what's happened so far, we know the history. And what is this history? It's going to tell us what happened in period 0 and what happened in period 1.

And I'm using the bar notation just to say this is fixed. So you could think of this as a, this is just a symbol. I'm writing the bar just to-- I think it will be a little clearer below what's going on, but this is a fixed action profile a 0, and this is a fixed action profile a 1. And as before, this tells me what every single player did in period 0, and this tells me what every single player did in period 1.

So now, we can think of our subgame as we're at a 0, we're at a 1. Here we are, this is our history. And then play is going to continue. The subgame is everything from here on out. And just to understand what's going to happen, let's say the way we play is going to induce some action profile a 2, some action profile a 3, some action profile a 4 and so on, and it's going to go all the way forward.

So let's try to understand what our payoff is going to be in this particular subgame. So our payoff is just going to be-- let's look at player i. The payoff for player i is going to be ui of a 0 plus delta ui of a 1 plus delta squared ui of a 2 plus delta ui of a 3, and so on.

So technically, if we want to compute our payoffs, we have to keep track of all the actions here. But a key observation is that this part isn't really going to matter for this subgame. Because once this is the history, and once we're considering how we're going to play moving forward, these actions have already been chosen. So these actions just serve as a constant.

Once I've made it to period 2, how I played in period 0 and period 1, it adds to my payoff, but it doesn't really matter for my decisions today because whatever I do moving forward, these have already happened and these numbers are already here.

And if we focus on this part, it's convenient to factor out the delta to the 2. So let's just factor this out. And we're going to get delta squared times ui a 2 plus delta ui a 3 plus dot-dot-dot.

And we know from Von Neumann-Morgenstern utilities is that if I take utility function and I add a constant to it, and I multiply it by a constant, that doesn't really change my preferences. So all that really matters is this. So often, and from here on out, when we're looking at a subgame, instead of writing out the payoffs like this, which is maybe formally correct, we're just going to directly jump to this because this is what matters.

And we're going to say, if you're in the period 3 subgame, all that matters is your payoff today plus delta times your payoff tomorrow and so on. And we only need to look at what happens moving forward, we don't have to look back at the past. So this is going to be-- this is what matters, and this is what we're going to use as the subgame payoff.

OK, so now that we've set that up, let's now try to actually construct an equilibrium. So let me come back over here. So let's consider the prisoner's dilemma. So remember, we have cooperate, defect, cooperate, defect. 2, 2, 3, negative 1, negative 1, 3, 0, 0.

So now our goal today, and for the rest of the class, is to find SPME of the infinitely repeated version of the prisoner's dilemma. So maybe I'll write PD of delta. What I mean here is the stage game is the prisoner's dilemma. We're going to repeat this game infinitely many times. And the discount factor we're going to use is delta. Both players use this discount factor.

So we want to try to find subgame perfect Nash equilibria of this game. And as context, let's recall, what did we see about PD of a fixed T? So the finitely repeated version of the prisoner's dilemma, we analyzed that on Tuesday. Does anyone recall what the subgame perfect Nash equilibria of that game were?

So remember on Tuesday we had a result that said, if a stage game has a unique Nash equilibrium, then the finitely repeated version of that game also has a unique subgame perfect Nash equilibrium. Yeah?

AUDIENCE: Because maybe they both defect.

IAN BALL: They just defect every period. And more importantly, they defect whatever's happened in the past. So the finitely repeated version of the prisoner's dilemma had a unique equilibrium, had unique subgame perfect Nash equilibrium. And what was it? It was si, s1 of h equals D for every history h. And similarly, s2 of h equals D for every h.

And again, we always have to be really clear about the difference between the outcome of a game and the strategies in the game. So the outcome of this subgame perfect Nash equilibrium was we defect every period. But just saying that we defect every period doesn't specify the strategies because the strategies have to say not only what we do do, but what we would have done had they been different.

So the strategies are this. We're saying not only do we actually defect, but we would have defected had the history been different. And in fact, here, then the subgame perfect Nash equilibrium is in every history, we always defect.

So now we want to understand how is the infinitely repeated version of this game different. We studied the PD of T, so how is the infinite horizon different? Any thoughts? Do people have intuition? Do you think people will be able to cooperate more? How do you the infinite horizon might be different? I don't know. Will this still be an equilibrium? Any thoughts? Or suggestions?

So it turns out that in the infinite horizon, we still have this subgame perfect Nash equilibrium, but we're going to have a lot of other equilibria as well. So the first point is that-- maybe I'll call it star. I'll call that strategy profile star. So star is still a subgame perfect Nash equilibrium of the infinitely repeated version.

Let's check this. So here's-- let's check. We have a strategy profile. We want to check that it's a subgame perfect Nash equilibrium. How do we do that? Well, we have to check that in every subgame, it gives us a Nash equilibrium. To make that easier, we're going to apply the one-shot deviation principle. So by the one-shot deviation principle, we only need to check one-shot deviations. One-shot deviations, but we do have to check them at every history.

I want to really emphasize this. So I think one common mistake that I see is people say, oh, we'll apply the one-shot deviation principle. We're looking for the equilibrium where we always defect, so let's just look at a history where we've always defected so far and then checked that no one wants to deviate.

But that's not enough. You have to consider every history, even histories that wouldn't happen under this strategy profile. So we have to take into account what happens at histories cooperate, cooperate, cooperate, even though those histories won't actually happen.

OK. So let's consider a fixed history h. And let's consider player i-- maybe player 1 to make it easier. So at some history h, let's consider player 1, they're choosing between two things. They can either choose defect or cooperate. So we're at history h.

Defect is what they're supposed to do because we want to check that always defecting is a subgame perfect Nash equilibrium. So we want to compare defecting, which is what they're supposed to do, to the one-shot deviation where instead of defecting today, they cooperate, and then from here on after, they return to their equilibrium strategy.

So what does the future look like if the player follows their equilibrium strategy? Well, today, they get DD. Why? Well, player 1 is defecting today, and the other player's strategy is to always defect. Then the next period, they get DD, The next period, DD, is your DD. Because this is just what happens in equilibrium. So if the player does what they're supposed to, and we follow the equilibrium path of play, then from today onward, we're just going to get defect-defect from here on out.

What if the player one-shot deviates to C? Now what happens? What does the future look like? What is the future path of play going to be? Any thoughts? Yeah, in front.

AUDIENCE: Should be the same, right, [INAUDIBLE]?

IAN BALL: So I guess what I'm saying is here-- so we're at history h in-- let me say in period t. And I guess maybe my notation was unclear. This is period t.

AUDIENCE: [INAUDIBLE]

IAN BALL: Exactly. So I'm including what happens today. So to go to this example here, if t equals 2 I'm starting right here inside the game. Yeah. Good clarification. So indeed, it'll be CD today. Why is it CD today? Well, my opponent is playing D today because their strategy says they always play D whatever the history. And I'm player 1, I'm deviating to C. We have to be clear, defect and deviate are not the same. So I'm deviating to the strategy C.

I play C, my opponent plays D, but now what happens up here? Now you said DD, I agree, but why? Why is it DD up here? Yeah?

AUDIENCE: It's only a one-shot.

IAN BALL: Because it's a one-shot deviation. We know that the second player is always playing D. So let's understand the reasoning for this. We know the future has to look like this for the second player because this is a unilateral deviation. So if it's a unilateral deviation by player 1, player 2 is simply going to follow their equilibrium strategy, and their strategy is to always defect, so we know that player 2 is always defecting.

What's tricky, though, is I could consider a more complicated deviation as player 1 where I cooperate today, and then I continue doing weird things in the future. But since this is only a one-shot deviation, once I get here, I'm supposed to do exactly what my strategy specifies, which is D. So we get DD, DD, DD.

And now, is this deviation-- this one-shot deviation profitable? We have to ask, is this future better than this future or not? Well, these look all the same, so the only difference comes down to this period. But if I do what I'm supposed to do and I defect, I get a payoff of 0.

But if I cooperate and the other player defects, I get a payoff of negative 1. And thereafter, the payoffs are the same, so this is better than this, and indeed, I've shown that the player does not have a profitable deviation. So we've checked that this is a subgame perfect Nash equilibrium.

So this is the equilibrium that looks a lot like the equilibrium of the finitely repeated game. The difference with the infinitely repeated game, however, is that there are also other equilibria where we do cooperate. So now the question is, how can we sustain cooperation in a subgame perfect Nash equilibrium?

So we could get into the math, but let's just think about it. You played this game with your friends yesterday or on Tuesday, and some of you were able to sustain cooperation. How did you do it? What did you think would happen. So if you played this game with your friends, and those of you who did play cooperate each period, why did you do that? What did you think would happen if you didn't cooperate? What was your strategy that you were using? What was your plan? Yeah/

AUDIENCE: --in this round if my friend cooperates. If my friend didn't cooperate, side effect, the next round, should always turn to defect, and then we would go from getting 2, 2 to 0, 0.

IAN BALL: Right. So generally the idea is, if my friend defects and chooses D and messes my payoff up, I'm going to punish them in the future by playing D as well. That's something called tit for tat, and we're going to analyze that a bit later. We're going to start with something even simpler, which is called the grim trigger strategy.

You can imagine very complicated things where maybe if my opponent defects, I'm going to punish them for three periods and then go back to cooperating, but here, we're going to think of a very vindictive player. And the idea of the grim trigger strategy is, if anyone ever defects, I'm going to defect from here on out.

So let's write this down. So we have s1 of h. We'll look at player one. It's the same-- oh, we can do i, It doesn't matter. So we're defining a strategy. That means we have to define what this player does at every history. And we're going to define this piecewise. We're going to say, well, there's some histories where no one has defected so far. So we're going to break up the histories in two.

We're going to say if h does not contain D, so what do I mean by this? A history that does not contain D means at this point in time, if we look back at the way the game has been played, no one ever defected. And that either means we're in period 0, so maybe no one defected because we didn't play at all, we're in period 0, or it means we're after period 0, but when we look back, people cooperated the whole time.

And then we have other histories where someone has already defected. We're looking back and some player has defected. So what's the idea here? If it doesn't contain D, and we've cooperated so far, intuitively, what should we do? We should cooperate, right? So far, everyone's cooperated, so we're going to cooperate. But if someone has defected in the past, then we're going to defect as well. So this is the grim trigger strategy.

The trigger is someone defecting. That's what triggers us to move from cooperating to defecting. It's grim because once someone defects, we defect forever after. This is the harshest, grimmest possible punishment we can impose.

I want to point out one strange thing about this strategy. When I say contains D, this means by either player. So in particular, let's say in period 0, the players play CD. Then what is the grim trigger strategy going to specify in period 1? Yeah? DD, right?

Now, for player 1, I think this is very natural. Let's think of this from the perspective of player one. They look back to period 0. They say, I cooperated, but my opponent defected, I'm unhappy with them, I'm going to punish them by defecting tomorrow-- or defecting today, sorry. We're in period 1, I'm going to defect today because yesterday, my opponent, player 2, defected. That makes sense. What about player 2, though? This is a little weird. What is player 2's thought process here? Yeah, in the front.

AUDIENCE: Maybe you're anticipating their punishment, so you're trying to limit as much as possible.

IAN BALL: Exactly. So what's weird about player 2 is they're looking back and they're saying, wait, player 1 cooperated with me yesterday. That's great. I'm the one who defected. I'm the bad guy. Yet nevertheless, under this strategy, player 2 is defecting today exactly for the reason you pointed out. I'm not defecting to punish myself, I'm defecting because I see that because I defected, I anticipate that my opponent is going to defect today, and if they're going to defect today, why don't I just defect as well?

So it has this weird feature that it might appear that players are punishing themselves, their own misbehavior, but really, they're anticipating the other player's punishment and responding to that, exactly as you said. Great observation.

OK, so now, we want to check, is this a subgame perfect Nash equilibrium? Well, the answer is going to depend on delta, and let's, first, I always think it's good in these questions to try to intuit which way it's going to go before you solve the algebra, it's good to build intuition, and it's also helpful on exams if you have a sense of what the solution should look like, and then you make an algebra mistake, you're able to catch it.

So it's going to depend on delta. Is your intuition that this is going to be a subgame perfect Nash equilibrium when delta is high or when delta is low? So I'm going to tell you, for some delta, this will be a subgame perfect Nash equilibrium. For some delta, it won't. It's going to depend on delta. For which one do you think it will work, high delta or low delta? Yeah?

AUDIENCE: You just guess what the higher the delta-- or if it's a high delta, then it will be a subgame perfect--

IAN BALL: I agree. And what's your intuition for that? Yeah?

AUDIENCE: Because you're getting a higher payoff really early, which means if you have a high delta, only the high-- only the earliest payoffs will [INAUDIBLE].

IAN BALL: So, almost. I mean, we have to think with the high delta, it's always the case that later payoffs matter less. But with a high delta, they just matter less-- the amount they matter is less.

So it's still true that the future matters-- the present is the most important period, but it's almost as important as-- sorry, the present is always more important-- today is always more important than tomorrow, but when delta is close to 1, tomorrow is almost as important as today is a better way of saying it. So I think that's part of the intuition.

And the other intuition, why do you think high delta is going to make this-- let's think of the incentives of the players. Let's suppose the player thinks about deviating-- or sorry, thinks about defecting. What happens if a player defects in this game under this strategy intuitively? Yeah?

AUDIENCE: They lose all future benefits of cooperating.

IAN BALL: Exactly, but what happens to them today? That's exactly right, and then today what happens?

AUDIENCE: They gain the extra amount.

IAN BALL: Exactly. So if we've cooperated so far, and I'm thinking about what to do, I face a trade-off. If I defect today, I increase my payoff today, but I get punished in the future by my opponent. So I'm trading off the benefits of defecting today against the losses I experience in the future. And if I'm patient enough, those losses in the future loom large and will discourage me from defecting, and therefore, we'll sustain a subgame perfect Nash equilibrium.

On the other hand, we can think of the extreme case, if delta is basically 0-- let's take delta to be 0. If delta is 0, then this is just the one-shot game. And we know in the one-shot game, the unique equilibrium is defect. So we know this can't be a subgame perfect Nash equilibrium when delta is 0. So intuitively, we expect, the higher delta is, the easier it will be to sustain this as an equilibrium, and let's check and confirm whether that intuition is correct.

I should make, I guess, one comment, just a technical comment. I'm applying the one-shot deviation principle. The one-shot deviation principle applies to multi-stage games that are continuous. It turns out, because we have this discount factor delta, these games are always going to be continuous, and they're going to satisfy that technical condition, and therefore, we can apply the one-shot deviation principle.

OK, so let's try to go through this. So I said when we're analyzing subgame perfect Nash equilibrium, the first step is always identifying the subgames, which is the histories in this case. I think the second step-- maybe I'll make this observation. The key step is grouping the histories correctly.

Because there's infinitely many histories, and we have to argue at every history, no player has a one-shot deviation. If we just went through every single history, we'd have no hope, I mean, we never finish.

So what we want to do is we want to split the histories into groups and argue about the entire group of histories, and then argue about the other group of histories. So what's the natural way you would group the histories here? If you look at the way the strategy is defined, I think it's natural to put the histories into two groups. What would those two groups of histories be if you wanted to reason about one class of histories altogether and another class of histories altogether? Yeah?

AUDIENCE: Yeah, like someone has defected or someone has--

IAN BALL: Exactly. So let's break the histories in two and let's consider-- maybe start with the easy case where someone has already defected. So maybe you can think of it as case 1. So we're proving something about every history. We're going to split the histories into groups. Maybe I'll say group 1 instead of case 1. Group 1. So let's consider history h that contains D.

So we're at some period, we're at some history containing-- contains-- containing D. And we need to check that at this history-- I mean, really, we're reasoning simultaneously about many, many histories, but we're going to-- let's just take one of them. And let's argue that neither player has a profitable one-shot deviation at this history. That's what we have to show. Everything here is symmetric. So I'm just going to reason about player 1, and all the reasoning I do would also apply to player 2. So let's focus on player 1.

And maybe a good way of saying this-- sometimes we say if I deviate-- so let's look at a one-shot deviation versus if I follow the strategy. Let's just understand what we're trying to show before we add all the details. We're considering a history that contains some D. So someone has defect in the past. We're trying to check that player 1 does not have a profitable one-shot deviation from this strategy, from this grim trigger strategy.

So we're going to compare what happens, if player 1 follows this strategy, what happens from here on out, versus what happens if player 1 chooses a one-shot deviation. In general, there could be many one-shot deviations because there could be many different actions the player could deviate to, but because this game only has two actions, there's only one one-shot deviation.

If I don't do what I'm supposed to today, there's only one other thing I could do. So instead of having many different one-shot deviations, namely the number of actions minus 1, we just have one of them.

And what I want to keep track of is what happens from today onward if I follow the strategy and what happens from today onward if I deviate. So let's start with what happens if I follow the strategy. What's going to be the future path of play from here on out?

Well, both of us are following this strategy, and the strategy says, if a history contains D, we better defect. So that means today, we're both going to defect. That's what happens if we follow the strategy. Again, defecting is an action. It's not the same as deviating from a strategy profile, so I'm just talking about the action defect. Then what happens the next day? Yeah?

AUDIENCE: Continue to defect?

IAN BALL: Continue to defect because if you've already defected here, and then we both defect here, well, certainly from this perspective, we look back, we see someone's defected, so we're just going to have D forever. What if player 1 chooses a one-shot deviation? Now what is the future going to look like? Let's start here. What's going to happen in this period if player 1 deviates? Well, player 1 is supposed to play D, so a deviation means they play C. So we're going to get CD today.

And then, from here on out, it's just a one-shot deviation. So everyone's going to follow the strategy. If we'd already defected at this point, certainly we've still defected here, so we're just going to have DD, DD, DD.

So, is this one-shot deviation profitable or not? No, because it's actually easy to see. Here, I do worse today, and then I do just as well from here on out. So certainly, this is actually strictly worse, so this is actually strictly unprofitable. So this one-shot deviation makes me strictly worse off than if I follow the equilibrium strategy profile.

OK, now let's look at group 2. Now I only reasoned about a single history and a single, but the reasoning is symmetric for player 2, and this reasoning applies to any history in the first group. So I kind of simultaneously dealt with many, many histories that look the same. Now let's look at group two. This is the harder case.

So group 2 is going to be all histories that don't contain D. So let's consider a history h that does not contain D. And it's a minor point, but why don't I say consider history h that just consists of C? You might think another way of saying it. Well, I'm making sure that I cover the null history.

So one history is, the game has started, nothing has happened so far. I want to make sure that counts, that history does not contain D, but it also isn't just C, so I have to say it this way.

And now let's go through the same reasoning. We have player one. And they can either follow the strategy, or OSD, One-Shot Deviate, and let's see what happens. So following the strategy is pretty clear. What's going to happen from here on out, if we know one has deviated so far, and we follow the grim trigger strategy, what does the future look like? What's the path of play? Yeah?

AUDIENCE: Cooperation every round.

IAN BALL: Cooperation every round. Because that's what we're doing. If no one's deviated, we cooperate. We cooperate again. No one's deviated, we cooperate again, and so on. What about the one-shot deviation? Yeah, over here.

AUDIENCE: I would defect in the first round.

IAN BALL: So let's break it up. Let's do the first round, yeah. So I defect here. And then what is-- and what's this going to be? Cooperate, OK. So that's the one-shot deviation, great. And then what happens after that?

AUDIENCE: I believe following the appropriate strategy after that, we would be defecting.

IAN BALL: Exactly right. So this is the key point and can confuse people a lot. I said this is a one-shot deviation, but what's happening is changing in every subsequent period. Can you reconcile that why? It's a one-shot deviation, but if I look at this, it doesn't look one-shot. So you are right, but this is a tricky point. Can you explain what's happening?

AUDIENCE: Well, so in fact, let's say time equals 0 right now, I'm going to make one decision.

IAN BALL: Right.

AUDIENCE: Switch my cooperate to defect.

IAN BALL: Exactly right.

AUDIENCE: Assuming this grim trigger strategy that we have set in stone before we started this game, based on that, the strategy dictates that every future period must be defect. And so based on that, then we are going to defect.

IAN BALL: Exactly, great explanation. So here, why am I-- why are we playing DD here rather than CC? It's not because a player is changing what their strategy does at this history, it's that we've reached a different history than before. So the function-- the strategy hasn't changed, but the history, what we're plugging into the function has changed, exactly as you explained.

So let's go over one last time. I one-shot deviate today. I'm supposed to play C, and I play D. Subsequently, I follow my strategy, but now, if we're here, and we look back, we say, wait a second, someone defected yesterday, the history contains D, and therefore, the grim trigger strategy specifies that we both play D, and we continue following this from here on out. I think this is the crucial step. Any questions on this?

OK. So now let's compare these payoffs. If we do CC every period, we get 2, 2, 2, 2. Now, what happens? I get-- so here, we can see exactly what's kind of the fundamental idea about punishments. If I deviate from cooperate to defect today, I gain today. We know I must gain today because CC is not a Nash equilibrium of the stage game. What it means for CC to not be a Nash equilibrium of the stage game is precisely that I can behave differently today and increase my payoff today from the stage game.

But I may not want to do that because if I do this, we're going to observe that tomorrow, and I'm going to experience a punishment from here on out. And in order for this not to be profitable, it must be that the punishment is sufficient to offset the gain, which is exactly going to be true if we're patient enough and the future looms large enough, so let's actually compute this.

That means we have to do a little bit of algebra here. Well, if we discount starting today, we're going to have discount this by delta, this by delta squared, this by delta cubed, and so on. So we're going to get 2 plus delta 2 plus delta squared 2 plus delta cubed 2 and so on.

Again, I'm using this trick that I'm focusing just on the payoffs here. It could be that the history we're at is period a billion, and everything is discounted by delta to the billion, but I've just factored that out to the front, and that's not going to change my preference, so I'm just going to remove that.

OK let's do a little simplification here. We see a 2 in every term, so let's factor that out. So this is 2 times 1 plus delta plus delta squared plus so on. And this is one series that I think, for this class, it's good to know the formula for this. This is a geometric series. It converges because delta is less than 1. And the formula for this is it's just 1 over 1 minus delta. So this is going to be 2 times 1 over 1 minus delta.

Let's just make sure this formula makes sense. If delta is really close to 0, then this series is basically just 1. And that makes sense, 1 over 1 minus basically 0 is 1. If delta gets very close to 1 this series gets really, really big. And indeed, as delta gets really close to 1, the thing in the bottom here gets really close to 0, and therefore, the reciprocal of it gets really, really big. That's Intuitive.

OK. Now let's compare that to what happens here. Well, this is actually pretty easy to calculate. What is this stream? Well, it's 3 plus delta times 0 plus delta squared times 0 plus dot-dot-dot. Well, the 0's don't matter, so we just get 3. So what is our equilibrium condition?

So what we see is that every history in this group, our deviation is not profitable as long as this is greater than or equal to this. And remember, equality is OK because if we have equality, then that means the deviation gives you exactly the same payoff as on path. Nash equilibrium allows that. So the condition is 2 over 1 minus delta must be greater than or equal to 3. Let's just do a little algebra here. Let's move the 1 minus delta to this side. We get 2 is greater than or equal to 3 times 1 minus delta, which is 3 minus 3 delta.

And now we move 3 delta over here to over here, and we get 3 delta greater than or equal to 1, or in other words, delta is greater than or equal to 1/3. It's very easy when you're doing this algebra to get the sign wrong, and that's why keeping track that we better get an inequality that has this form, delta greater than something because we know that this should be an equilibrium if and only if delta is large enough.

So if I made a mistake and I got delta less than or equal to 3, i would know I made a mistake if I understand the economic intuition, so that's a good tip for exams.

OK. So we intuited that this strategy would be a subgame perfect Nash equilibrium if the players are patient enough, if delta is large enough, and our calculation found the exact threshold, which is exactly 1/3. And that threshold reflects the ratio between the one-shot gain from defecting and the subsequent punishment that I experienced. OK, any questions on that? All right, let's do one more strategy. Let's see what happens. Yes, absolutely.

AUDIENCE: I just want to confirm that basically, the CC path only works if delta is large enough. So basically, CC is a SPME if and only if delta is greater than or equal to 1?

IAN BALL: We have to be really clear about what-- remember, an SPME is a strategy profile. So what we're saying is the strategy profile where both players use the grim trigger strategy is a subgame perfect Nash equilibrium if and only if delta is large enough.

Now you're right that at histories where someone has already defected, there's never a profitable deviation no matter the value of delta. If we go to group 1, this deviation was never profitable. But to be a Nash equilibrium, it's not enough to say there's some histories where the deviation is profitable. SPME is really strong. It says at every history, no deviation can be profitable. So that's only going to hold if delta is greater than or equal to 3.

So another way of saying it is, if delta is large, whatever history I'm at, no one has a profitable deviation, and therefore, I have a Nash equilibrium. If delta is smaller than 1/3, there are some histories where I have a profitable deviation and other histories where I don't have a profitable deviation. But as long as there's some history where I have a profitable deviation, that does not constitute a subgame perfect Nash equilibrium, and therefore, the strategy profile doesn't work as a subgame perfect Nash equilibrium.

AUDIENCE: And if the deviation is greater than 2?

IAN BALL: The profitable deviation. I mean, I have deviations here, they're just not profitable. But the profitable deviation will always be at group 2 histories, yeah. So in fact, there's-- keep in mind, though, this is when you call it group 2, it's not just one history, it's actually many, many histories that all fall into this category, yeah. Great. Any other questions on this? Yeah?

AUDIENCE: Given the Nash equilibrium profile, do we specify which one we end up in the beginning? Like, how do we know if we end up in DD or CC?

IAN BALL: Great. So, great question. So as always, we have to distinguish between the strategy profile and equilibrium and the outcome of the equilibrium. I've described the strategy profile, but I should say what the outcome is. So let's look at the outcome.

Well, the strategy profile says-- maybe it's gone now, if no one's defected so far, we cooperate. So in the 0th period, we should cooperate because no one-- it's the 0th period, no one's defected so far, so we cooperate. Now let's go to period 1, what do we do?

AUDIENCE: We continue--

IAN BALL: We cooperate because no one's defected so far. Period 2, we cooperate. So the outcome of this equilibrium is always cooperate, and maybe I should have said that. So the outcome-- I was so focused on strategies versus outcomes, I didn't even say what the outcome is. The outcome is CC, CC, CC, CC. So that's the outcome.

We do, indeed, sustain cooperation in equilibrium, but I didn't say this because I want to be clear, it is not correct to say the equilibrium is cooperate forever because that's an outcome, not a strategy profile, but great question. Yes?

AUDIENCE: Is this outcome still under the assumption that delta is at least 3?

IAN BALL: Yeah, yeah. So let's be clear-- well, I guess it depends what you mean. The outcome of the grim trigger strategy profile is always-- is C, C, C, C, C. But that strategy profile is only a Nash equilibrium if delta is greater than or equal to 1/3. For any strategy profile, I can compute what the outcome is, whether or not it's equilibrium. I can say, if this is how we-- if these are our complete contingent plans, this is what's going to happen. That doesn't rely on equilibrium.

But then a separate question is, are those plans actually consistent with equilibrium? And that's where we need delta greater than or equal to 3. Is that clear? Yeah. Any other questions? Great. These are great questions.

OK. So now let's look at one more, I think, very natural strategy profile and see if this works. So this is called tit for tat. And people actually have these contests where you play prisoner's dilemmas, and this strategy actually tends to do quite well. Tit for tat-- don't know the etymology exactly what tit and tat mean, but the basic idea is I'm always going to do what you did yesterday. So if yesterday you cooperated, then I'll cooperate today. If yesterday you defected, then I'll defect today.

So tit for tat means each player copies the opponent's action yesterday. So opp means opponent. If you cooperated yesterday, I cooperate today. If you defected yesterday, I defect today. And that seems like a pretty reasonable strategy.

Let's be careful. I need to say one more thing. I actually haven't fully specified a strategy here. What have I not specified? And this is crucial. Yeah?

AUDIENCE: What you do in period 0.

IAN BALL: What I do in period 0. In period 0, nothing's happened so far, and that makes a huge difference. So plus cooperate in period 0. So when we say tit for tat, that's usually what we mean. So in period 0, I cooperate. And in any subsequent period, I copy what everyone else did.

So the first question we should ask, going back to our discussion we had, is if we use the tit-for-tat strategy, what is the outcome going to be? And remember, we can ask this question whether or not tit for tat is a Nash equilibrium. We can just say what would happen if this is played. So what is the outcome of this going to be? We're not saying it's a Nash equilibrium or subgame perfect Nash. We're just saying if these are the contingency plans people use, what's going to happen? Yeah?

AUDIENCE: Because everyone cooperates in period 0, subsequently, everyone cooperates period 1, and then the outcome is that CC has played every round.

IAN BALL: Exactly. So you might say, great, this induces CC, this seems reasonable, maybe this is another equilibrium that's going to allow us to sustain cooperation. We already computed grim trigger, let's see if this one works. It turns out, it's not going to work, even though it seems quite intuitive, and let's understand why.

So what we want to do is we want to check, is-- maybe I'll say TFT, Tit For Tat-- a subgame perfect Nash equilibrium? Well, I need to check, at every history, does anyone have a profitable one-shot deviation? And as usual, there's a lot of histories. I can't reason about all of them, so I need to break up the histories into groups.

Last time we grouped them into two groups according to whether someone had deviated or defected in the past. How would you group the histories under this strategy? Any guesses how many groups you might have and how you would group them? Yeah?

AUDIENCE: Someone deviated yesterday versus someone cooperatinbg yesterday.

IAN BALL: Great. So that-- so I agree, we need to look exactly at what happened yesterday. And whether someone deviated tomorrow is important. It turns out, we want to keep track of all the four possibilities yesterday, and maybe this is what you meant. So yesterday, it could have been we both deviated-- sorry, defected. I always mix these up. We both defected, we both cooperated, or I cooperated and you defected, or you cooperated and I defected. So exactly right, but there's going to be four of them.

So we're going to group them into four groups based on what happened yesterday. So maybe I'll say group 1 is a history h with CC yesterday. And now we're going to have CD, DC, DD. OK, great. Yes?

AUDIENCE: Just out of curiosity, the question, how is DD possible? How is that one-shot deviation to DD?

IAN BALL: Great, but remember, the crucial thing about the one-shot deviation principle is we have to consider every history, even histories that aren't possible. So you're exactly right, not all of these histories are possible. Given the strategy, and this is a common thing, but the good news is, this makes your life easier.

You don't have to figure out at which histories are possible and which are not. Don't worry about it. Just you have to look at every history, and in fact-- yeah, I think this is the most counterintuitive thing about all this, but that's a crucial thing. At every history, even histories that aren't reached. Why? Well, because of non-credible threats.

It may be that the reason history was not reached was that we were doing something ridiculous at that history, and therefore, we still have to figure out how people are playing at that history. That's the key point here. Great.

Again, I need to be a little careful. I'm missing one thing. What's the gap here? I have-- what history's not covered? The null history. There's no yesterday. But since we cooperate in period 0, let's group the null history with this group. So history-- with CC yesterday, or in this case h equals the empty set or with CC yesterday. So we'll group this history into here.

OK, so now every history is in one of the four groups. Mathematically it's called a partition. We've split up the histories into groups. And now we want to go through our same reasoning here. Again, let's just look at player 1. Everything symmetric. So we could look at player 2, we'd get the same answer, so let's look at player 1. And we want to compare what happens if they follow versus a one-shot deviation. So maybe I'll say follow or one-shot deviation, and then we'll go through all these.

So, the history is CC. I think for this, it's easier to go step by step. So let's fill this in. Yesterday, we played CC. So today, if we follow the equilibrium strategy, what happens? CC, and it's going to continue that way. What if I do a one-shot deviation? I would encourage you to answer this in two steps.

So first, let's say what happens today. If I deviate today, well, the outcome is DC, because my opponent was playing C today anyway. I was supposed to play C, but I'm deviating, so I'm playing D.

Now that we figure out what happens today, let's figure out what happens tomorrow. It's a one-shot deviation. So from tomorrow onward, we're both following the equilibrium strategy of tit for tat. So if DC was played today, what's going to be played the next period? Yeah? CD, right? Why? Well, player 1 is copying what player two did, and player 2 is copying what player one did.

And then you can see, we're going to keep alternating DC, CD, and so on. There's a lot of cases-- we're running out of time, let's just do case 2, and we'll make an observation. So follow OSD. Actually, Let's skip case 2 and go to case 3. It's going to be more helpful. I encourage you to do cases 2 and 4 on your own, but this is going to be a more useful case.

So if we follow, what happens? Well, yesterday was DC, so today is going to be CD. Player 1 copies what player 2 did and player 2 copies what player 1 did, and then we're going to see the same pattern. Maybe I should have done yesterday. OK. Anyway, this is-- OK.

And if I have a one shot deviation as player 1-- I did this wrong-- we get-- my original ordering was right. Sorry. Let's go-- let's go-- OK. OK, let's look at the case CD. Sorry, this is the more important one. If CD was played yesterday and I follow, we're going to get CD, DC, CD, and so on.

AUDIENCE: We'll get DC first, right? And CD was yesterday, then.

IAN BALL: Ah, yes. Thank you. Thank you. And-- exactly. Yeah, I'm confusing myself here. And what if we one-shot deviate-- ah yes. Anyway, OK. What if I one-shot-deviate today? I'm supposed to play DC. I'm player 1, I defect-- or I deviate to CC, and then we're going to continue like this.

I see. So let me actually-- let me erase this one because it's just causing confusion. It isn't actually the important one. I switched things. OK. Let's go over this slowly.

We played CD yesterday. We're supposed to play DC today. And if we continue, we're going to get this. If we were supposed to play DC today, but I deviate, then instead of playing D, I play C, I'm player 1, but if we both cooperate today, then from here on out, we get cooperate. OK.

Now we need to check that neither of these is profitable, but I argue that if we look really carefully, we can see that this is not going to work in general. What's the issue? Yeah?

AUDIENCE: [INAUDIBLE]. CC is better than all of these.

IAN BALL: Ah, but DC is actually better than CC, so it's not obvious. Yeah?

AUDIENCE: If we say one of these patterns and say DC, CD, so on is more profitable than CC overall, so the one-shot deviation for case 1 is profitable. Then by default, the one-shot deviation for case 2 cannot be an improvement because the exact reverse--

IAN BALL: Exactly. So the key observation is, what we get if we follow here is what we get if we deviate here, and what we get if we deviate here is what we get if we follow here. So if the stream DC, CD, DC is higher, is better, then we have a profitable deviation in this case. If the stream CC, CC is higher, then we have a profitable deviation in this case.

So in fact, the only way neither can be profitable is if these two streams are exactly the same. And I won't do this, but you can actually check that this holds if and only if delta is exactly 1/3. So it's kind of a degenerate case. For most deltas, this is not going to be equilibrium, and I think for this reason, we would mostly say we don't think tit for tat is a very compelling prediction because you have a really knife edge case. If delta is slightly higher or lower, it's not going to work.

You can go through this case as well and check that, indeed, if delta is 1/3, this will be a subgame perfect Nash equilibrium, but for any other value of delta, which is basically always the case we're in, this will not be a subgame perfect Nash equilibrium. Let me stop there. Great.