[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: So today we're going to talk about Nash equilibrium. So this is one of the most important concepts in game theory. But I want to start by reminding you a little bit about rationalizability because I think that's going to provide us a good foundation for talking about Nash equilibrium, seeing how these two concepts differ.

So as before and throughout this class, we're going to be given a strategic form game. And this notation is going to stay the same. So we have a strategic form game. Maybe I'll write strat-form game, which specifies the set of strategies, the n players, and then the payoffs that each player gets as a function of the strategy profiles that are chosen.

And last time, we defined this procedure, which we called iterated elimination of strictly dominated strategies. And this was an iterative algorithm that gradually narrowed down or pared down the strategies that the players had. So we started with, if you remember, the first stage or stage 0, the initialization stage was just the set of all strategies, all strategy profiles.

And then we went to big S, S1, S2, and we kept going down until maybe we had a generic SK, and then, at the end, we had something called S infinity. And it's important to understand-- maybe I'll write this up here-- the interpretation of each of these sets. So below, let me say what is the interpretation?

The interpretation is that, as we're willing to make stronger and stronger behavioral assumptions about the way that the players play, we're able to make sharper and sharper predictions about the strategies that we could observe in practice. So we start with S0. This imposed no assumptions.

If we make no assumptions about how people play, then, of course, anything is possible. The set of strategy profiles that the players may choose is unrestricted, but then we move to S1, where, there, the interpretation is rationality.

What we said is, what if we assume that every player is rational in the sense that every player plays a best response to some belief about the opposing player's strategies? And if we assume that every player is rational, then we can narrow down our prediction in the game to the set S1. So capital S1 is the set of strategy profiles that are consistent with every player's rationality.

And just to remember-- maybe I'll put it down here-- when I write S capital K, this is really a product. Like this. So S capital K is a set of strategy profiles that contain a strategy for player 1, in this set S1 K, all the way up to a strategy for player n in the set SN K.

And then we could go beyond just rationality. Remember, we went to S2, and there, we assumed not only that the players are rational but that the players know that everyone else is rational. So we might call this rationality plus knowledge of rationality. I'll just write knowledge of R, so rationality plus knowledge of rationality.

But then, as we keep going down, we can get to higher- and higher-order knowledge of rationality. So when we get to SK, the interpretation of this is SK is our prediction in the game if we're willing to assume rationality, knowledge of rationality, knowledge of others' rationality, knowledge of others' knowledge of others' rationality, and so on, all the way down to maybe I'll say rationality plus "K minus 1 order" knowledge of rationality.

Why K minus 1? Well, the first stage, we had rationality but no knowledge of rationality, so the degree of knowledge of rationality is always just 1 smaller than our number K. And when we get up to K, this means the players are rational. They know everyone else is rational, they know that everyone else knows that they're rational, and so on. And those statements go on until I say the word knows k minus 1 times.

And then if we want to be really formal, there's a word for this last set. So this last set, these are the set of rationalizable strategies. And the interpretation of this is that these are the strategy profiles that are consistent with rationality and arbitrarily high degrees of knowledge of rationality. So not only does each player everyone else knows their rational, but I can continue those chains arbitrarily deep. And we have a word for that-- that's called common knowledge of rationality.

So I'll write-- just write R. So rationalizable captures the idea that not only are the players rational, but they have common knowledge of rationality in the sense that they have K-th order knowledge of rationality for any K. They have first order, second order, third order, fourth order, 100th order, and so on.

So this is kind of a summary of where we've gone so far. And I want you to notice how, as we continue to the right here, we make stronger and stronger behavioral assumptions, and we get sharper and sharper predictions. But sometimes, even this set S infinity is not going to give us a very sharp prediction. And we want to go a step further. And that's what Nash equilibrium is going to do today.

So to motivate that, let's start with or revisit an example that we've done before, which we call the Boston game, BOS. And remember, the story of this game is we had two friends who like each other's company but have different sports preferences. So they're choosing whether to go to a Celtics game or a Red Sox game. They're both happier if they go together than apart, but one person prefers the Celtics and one person prefers the Red Sox.

So if we write out the game like this-- I think I'll use slightly different numbers this time to make some of the math later a bit easier. So let's go over it. If we choose differently, that's along this off diagonal, then we both get a payoff of 0 because we don't get to be together at the game.

If we both go to the Celtics game, then we both get a positive payoff. But it's higher for player 1 because player 1 likes the Celtics more. And then conversely, if we both go to the Red Sox game, we both get a positive payout. But here, it's the second player who prefers the Red Sox and therefore, gets the higher payoff of 2.

So let's go through our procedure for calculating best responses here. If my opponent is going to the Celtics game, I certainly want to go to the Celtics game. And if they're going to the Red Sox game, I want to go to the Red Sox game. And then we can do it the other way.

Now, I'll put myself in the shoes of player 2. I'm saying, if player 1 plays C, well, I look at my payoffs. I prefer C. If player one plays R, I look at my payoffs, I prefer this.

So now, if we look at that, can anyone see, what are the set of rationalizable strategies or strategy profiles in this game? So maybe we'll separate it out just to be clear. So what are the set of rationalizable strategies for player 1 in this game?

Well, here, it's going to be just C And R because, here, if my opponent plays C, I want to play C. If my opponent plays R, I want to play R. So it turns out everything is going to be rationalizable in this game.

So we're going to get-- if we go through the procedure C, R, and if we go through the procedure, we also get C, R. And just to make sure we understand the notation, S infinity is the set of rationalizable profiles. So this is going to have four elements. It's going to have CC, CR, RC, and RR. Sometimes I'll write with parentheses, but it just gets a little messy here.

So this says there's four strategy profiles. This is the profile where we both go to the Celtics game. This is the profile where we both go to the Red Sox game. And the order here indicates who does what. So the first player plays C, the second player plays R, and then, conversely, the first player plays R and the second player plays C.

So in this case, everything is rationalizable. So this is a case where rationality and common knowledge of rationality alone doesn't give us a very good prediction or doesn't rule out many possibilities. So we're not able to get a very sharp prediction.

And let's see, so in particular, one rationalizable strategy, profile, I should say, is this. So this is rationalizable. But it may seem like something is going wrong here.

So does this seem like a reasonable prediction of how people are playing? Or would you say someone's making a mistake here? What seems strange about this-- when player 1 is going to the Red Sox and player 2 is going to the Celtics game? Yeah.

AUDIENCE: [INAUDIBLE] since they're both losing.

IAN BALL: They're both losing. And in a sense, well, I'm going to the Red Sox game, but if I just changed my behavior and went to the Celtics game, I'd be better off. And similarly, the other person who's alone at the Celtics game, if they move to the Red Sox game, they'd be better off.

So there's a bit of this incongruity that we say, OK, this is consistent with rationality and common knowledge of rationality, but it might feel somehow that the players are making a mistake. So maybe I'll say this is all a bit informal. It feels like a mistake.

And let's try to reconcile this. Why does this satisfy the definition of rationalizability but feel like a mistake? And I think the best way to think about that is, what is rationalizability really mean?

It really means that a player has an explanation for why they're doing what they're doing. So I think rationalizability should be thought of as it means that a player has an explanation. So what is player 1's explanation for why they're going to the Red Sox game?

They would say, well, it's reasonable, it's rational for me to go to the Red Sox game because I believe that my friend is also going to the Red Sox game. And then we can go a step deeper. We could say, well, why do you believe that your friend is going to the Red Sox game? And you could say, well, I believe my friend is going to the Red Sox game because I believe that my friend believes that I'm going to the Red Sox game.

And then you go a step deeper. You can say, why do you believe that your friend believes that you're going to the Red Sox game? Well, I think my friend believes that I believe that my friend believes he's going to the Red Sox game, and that explains his belief, and we can go on and on and on.

And that's precisely what rationalizability captures. It says you can provide an explanation that goes arbitrarily deep. So think of this as the answer to the child who always says, but why this?

You can always go a step deeper. Why do they have this belief? Well, they have this belief because of this other belief.

And that chain never stops. If it weren't rationalizable, eventually, we'd reach a problem. We'd reach a contradiction in that chain.

So rationalizability says, there exists some explanation for your play. The problem is your explanation may not be correct. In theory, it's a good explanation, but it might not actually comport with the way that people are playing.

So rationalizability means you have an explanation. You're behaving optimally, given this explanation or given these beliefs, but there's no requirement that the explanation is correct, that your beliefs or your explanation-- I'll use these explanations a bit deeper. Because if we're down here, my explanation involved my belief that my friend was also going to the Red Sox game, but in fact, my friend is not doing that. My friend is going to the Celtics game, and that's why I'm doing badly.

So what Nash equilibrium is going to capture is it's going to go a step beyond simply providing an explanation. It's going to require that the explanations we provide are also correct. That is, they also accurately describe the way that our opponents or that our fellow players are actually playing.

So let's now give the formal definition of this. So if you learn one thing about game theory, this is often the definition you see. In the movie about Nash, they talk about this.

Though, actually, their story is kind of wrong in the movie. It doesn't make mathematical sense. But anyway.

OK, so a strategy profile-- let me just stop here and say, I think this is really an important aspect of the definition, that the thing that a Nash equilibrium will be is a strategy profile. And often, a mistake students make on exams is they're asked to give an equilibrium, and they often don't give a strategy profile. They might give the payoffs that the players get if they use that strategy profile.

But an equilibrium is always a strategy profile because it's a prediction or recommendation for how people will play in the game. So this is the key first part. It's a profile specifying a strategy for the first player, all the way up to a strategy for the n-th player.

So this kind of strategy profile is a Nash equilibrium. So this is named after John Nash who won the Nobel Prize in economics. He wrote this paper in '48, '49. He wrote a few papers that introduced this concept.

It's actually interesting to read the original paper. The original paper is like two paragraphs long, and he got the Nobel Prize for it. So if you have good ideas, you don't need to go on and on and on. So a strategy profile is a Nash equilibrium if, for every player i, the following holds.

So I'm going to put myself in the position of player i, and I'm going to compare the payoff I get from my strategy to the payoff I could get if I deviated to a different strategy. So I'm going to compare ui of si star on the left to ui of si prime on the right. So I'm comparing-- I'm player i, and I'm saying if I choose si star, what I'm supposed to do in equilibrium, I want to do better weekly than if I choose some other strategy si prime.

But I still have to fill this in. And this is going to say, given the way that the other players are playing. So what I'm going to fill in here is s negative i star and s negative i star.

I can't control how the other people are playing. So I take that as fixed. I take that as given.

Given The way that my opponents are playing, I just do what I can control, I compare the strategy I'm supposed to choose to any deviation, and I require that the strategy I'm supposed to choose is weakly better for all si prime in si. In particular, si prime could just be si star, but then this is going to be satisfied automatically because both sides are going to be the same.

So let's think of another way to say this. Another way to say this-- so this is equivalent to this statement here-- is that si star is a best response to s negative i star. So another way of defining a Nash equilibrium is to say that every player I look at, if I look at the strategy that player's using, that strategy better be a best response to the strategies that the other players are using.

Now, we have to be a bit careful here because, when we first defined a best response, we defined a best response to a belief, not to a strategy of the opponents. But when I say a best response to this strategy profile, what I really mean, if I want to be really formal, is I mean to the belief that assigns probability 1 to my opponents playing this strategy profile, so to, more formally, the belief beta negative I with beta negative i of s negative i star equal 1.

So if my belief is my opponents are certain to play s negative i star, then the best thing for me to do or a best thing for me to do is to play si star. And here, you can immediately see in the definition the consistency that we're imposing on beliefs. So I think keeping in mind this interpretation is crucial because, when you first see the definition, you might say, where are the beliefs?

Well, because we're requiring consistency of beliefs, the beliefs are the same as what's actually happening. So s negative i star is how the other players are actually playing. And beta negative i is my correct belief that the players are playing exactly how they are playing. So sometimes we can lose track of this distinction between play and beliefs because our definition requires that they agree.

Maybe we can give a bit of a less mathematical frame on this. So another way we might say it is we might say that no player has a strictly profitable unilateral deviation. So let's break down what this means.

No player-- so let's focus on a single player-- should have a strictly profitable unilateral deviation. So a deviation for a player is a strategy that deviates from or is different from their conjectured equilibrium strategy. So we start with si star for player i, and we contemplate an alternative strategy for player i.

A unilateral deviation means I'm the only one who's changing my strategy. And this is I think the key, maybe the most important idea that Nash introduced was that, when I'm reasoning about other people, I should take their behavior and their strategies as fixed because I can't directly control what they do. So when I contemplate deviating, I only consider a change in my strategy.

And you can see that up here, in this inequality, that the way that my opponents are playing is the same on both sides of the inequality. The only difference between the two sides is how I'm playing as player i. And this strategy cannot be-- this deviation cannot be strictly profitable.

It could be that you have a deviation that gives you the same payoff as your equilibrium strategy, but it can't be that you can deviate and do strictly better. And again, we're using profitability not in the sense of a monetary payoff but in the sense of utilities. So a profitable deviation is a deviation that gives me a higher utility, a strictly higher expected utility.

So let's go back to this game and see if we can find the Nash equilibria of this game. Any guesses on what the Nash equilibria of BOS this game are going to be based on the definition or based on the things we've written? Yeah.

AUDIENCE: Both [INAUDIBLE]

IAN BALL: Right, so these are going to be the two Nash equilibria. How can we see this? Well, let's suppose we're both going to the Celtics game. Does anyone have a profitable unilateral deviation?

No, because if I were to unilaterally deviate, we're together at the Celtics game, if I deviate, I'm by myself with the other game, so I'm certainly worse off. And the same is true if we go to the Red Sox game. If we're both going to the Red Sox game, if I unilaterally deviate, I'm going to be worse off.

Now, of course, if we're both going to the Red Sox game and I'm player 1, I would much rather that we both go to the Celtics game because I'm a Celtics fan. But moving from here to here is not a unilateral deviation. All I can control is whether I go to the Red Sox game or the Celtics game.

Of course, I would love to force you to come to the Celtics game with me, but that's not something I can do in this game. So that's the intuitive explanation. Mathematically, we can see that both numbers are underlined.

And that's actually a great way to compute Nash equilibria in games like this. So we go through for player 1. We underline all of their best responses. Then we go through player 2 and underline all of their best responses. And we look-- so I guess this is an algorithm for finding them.

So step one is underline best responses. And it's crucial that you underline both players best responses and underline all best responses. So let me say all best responses for both players.

And then step 2 is you just check which cells-- or which boxes in the matrix have both numbers underlined. So look for cells-- what is a cell? Well, a cell in this context is just a strategy profile. So let's remember that.

And if you think through this carefully, you'll see this is just another way of saying the definition. If both cells are underlined-- well, if the first player's number is underlined, that precisely means player 1 is playing a best response to player 2's strategy. And if player two's number is underlined, that means player 2 is playing a best response to player 1's strategy. And that's exactly our definition of a Nash equilibrium, at least in a two-player game.

So this is a good way going through and finding and checking for Nash equilibrium. So now I think it's helpful to try to understand the relationship between some of the solution concepts and some of the definitions we've given so far. So so far, we have three, I think, important definitions.

One was rationalizability. And we represented the set of rationalizable strategy profiles. We called that s infinity.

Then another important notion is Nash Equilibrium. And let's call this NE, so let me think of NE as the set of Nash Equilibrium strategy profiles in a game. So just to be clear, this is a subset of S.

What kind of objects are these? These are strategy profiles. So rationalizable strategy profiles live in the set of all strategy profiles. Then the set of Nash equilibria that's also a subset of big S because a Nash equilibrium is also a strategy profile.

And then we had another notion, remember, which was Dominant-Strategy Equilibrium. And we'll call that DSE. And that's also a subset of S. So we have three sets of strategy profiles here.

Let's remember a dominant-strategy equilibrium, if you recall from the dominance lecture, was a strategy profile where each player's strategy is a weakly dominant strategy. That was the definition. So now what we'd like to understand is the relationship between them.

If something's a dominant strategy equilibrium, is it a Nash equilibrium? If it's rationalizable, is it a Nash equilibrium? And so on. I think a good way to get started, when you're approaching a question like this, is to work through an example and see if you can get some inspiration from that.

So let's look at our BOS example. So here, the set of rationalizable strategies was all four strategies. The set of Nash equilibria in this game over here with just two strategy profiles right, CC and RR.

What about the set of dominant-strategy equilibria? Which profiles were dominant-strategy equilibria in this game? Any thoughts?

Well, let's go through for player 1. Let's check. Is C a dominant strategy for player 1? That is, does C always do better than R and sometimes do strictly better?

Well, no, because if the other player plays R, then C does worse than R. C gives me 0, and R gives me 1. So C does not weakly dominate R.

And now let's ask the other question. Does R weakly dominate C? Well, no, because if my opponent goes to the Celtics game, I'd rather go to the Celtics game than go to the Red Sox game.

So neither does C dominate R, nor does R dominate C. And that means that neither of them is weakly dominant because if it was weakly dominant, it would have to weakly dominate everything else. And therefore, we have no dominant strategy equilibria in this game. So here, we have none. Or we might say we have the empty set.

So at least in this example, we see that Nash equilibrium is more demanding than rationalizability. The Nash equilibria are a subset of the rationalizable strategy profiles, and dominant-strategy equilibrium is even more demanding still. It's a bit weird to say it in this case because there are none of them, but it's vacuously true that every dominant strategy equilibrium is a Nash equilibrium. They just aren't any of them. But if there were one, it would be a Nash equilibrium.

So that gives us maybe a conjecture based on this example about the relationship between them. And I chose the example, so we wouldn't be led astray. And indeed, this is a result that, in general, we have DSE subset of NE subset of S infinity, and maybe I'll write subset of S just to remind us the space that we're in over here.

So let's understand what these subset relationships mean. DSE is a set. It's the set of all dominant-strategy equilibria which is a set of strategy profiles. This is another set, and this is another set.

So this inclusion tells me every dominant-strategy equilibrium is necessarily a Nash equilibrium. The converse is not necessarily true. It could be that something's a Nash equilibrium but not a dominant-strategy equilibrium. Can you come up with an example of a Nash equilibrium that's not a dominant-strategy equilibrium, maybe in this game? Yeah.

AUDIENCE: [INAUDIBLE] is a Nash equilibrium, but it's not a dominant strategy.

IAN BALL: Right, and same with R, so here, it's easy because there are no dominant-strategy equilibria. So both of the Nash equilibria are examples of Nash equilibria that are not dominant-strategy equilibria. And then, similarly, we have, there are some rationalizable strategy profiles that are not Nash equilibrium.

In this example, CR and RC are two profiles that are rationalizable but not Nash equilibrium. Now, sometimes these sets could agree, but sometimes this set could be strictly bigger. But our claim here is just that everything in here is in here, and everything in here is in here. And now we'd like to go through the reasoning for that, but are there any questions on what I mean by this statement before we go through the argument?

OK, so let's try to think through it. So here's the argument. There's two different things we want to show. I mean, let's be clear, this is just for context.

This relationship is obvious. It's just, by definition, every rationalizable strategy profile is a strategy profile. So we have two things to check. So let's do the first one.

The first one is that DSE is a subset of NE. So that's a mathy way of saying it, but what we mean is if you give me a dominant-strategy equilibrium of some game, I know that it must be a Nash equilibrium of that game. So let's see how we can show that.

So first, let's start with a dominant-strategy equilibrium. So suppose S star, which I'll write S1 star up to SN star, is a DSE. That's a starting point.

We're just going to take a strategy profile and make the assumption that it is a dominant-strategy equilibrium. And we want to now reason through why this strategy profile must also be a Nash equilibrium. Well, to show that it's a Nash equilibrium, we have to check that no player has a profitable unilateral deviation.

So let's consider a player. So that means that for, each player i, let's consider player i, and let's also consider a potential deviation by player i. So for each player i and each strategy si prime, what do we want to show?

We want to show that this deviation is not profitable for player i. That is, we want to show that, for player i, it's better for them to play si star than to play si prime, given the way that their opponents are playing. But that's what we want to show.

All we know is that this is a dominant-strategy equilibrium. So what's something useful-- do you have any ideas of how we could use the fact that this is a dominant-strategy equilibrium to get us started? What do we know about si star, given that s star is a dominant-strategy equilibrium? Yeah.

AUDIENCE: si star [INAUDIBLE]?

IAN BALL: Yeah, so let's go through it in turn. Because it's a dominant-strategy equilibrium, we know that si star is weakly dominant. And what does it mean to be weakly dominant?

It means that si star weakly dominates every other strategy of player i. But in particular, that means si star weakly dominates si prime because that's another strategy of player i. So we know that si star weakly dominates si prime.

We know more than that. But here, we're just throwing away information. If si star weakly dominates everything, in particular, it weakly dominates si prime.

And now can we reach the conclusion we want? Well, yeah, if si star weakly dominates si prime, that means si star always does weakly better than si prime, no matter how my opponents are playing, so if that's true and sometimes strictly better. So if that's true no matter how my opponents are playing, it's certainly true in the particular way my opponents are playing now.

So that tells me that ui of si star si negative-- s negative i star is greater than or equal to ui si prime s negative i star. So because si star weekly dominance si prime, this inequality would be true whatever I put on both sides here. I could put any as long as I put the same thing on both sides.

So in particular, if I put the way my opponents are actually playing, this is going to be true. And this precisely means that si star is a best response to s negative i star. So I conclude si star is a best response to s negative i star?

Well, I have reason for player i, but I can go through the exact same argument for all the other players. So not only is player i playing a best response to the way everyone else is playing, but every single player i is playing a best response to the way everyone else is playing. So I'll say, therefore, s star is a Nash equilibrium.

So if you're not so familiar with writing proofs, this kind of reasoning may be a little tricky. Let me pause there. Are any questions on this?

So it's really just about being organized and going step by step. None of the steps is too difficult, but you have to keep in mind where you're going, where you're starting, and carefully go through the logic of it. But the basic idea is pretty simple. If what I'm doing is best no matter how my opponents play, it's certainly best, given the particular way that my opponents are playing. And that's the high-level idea of this.

OK, so now let's try to show the second part. So I have shown that every dominant strategy equilibrium is a Nash equilibrium. Notice, I didn't show that there exists a dominant-strategy equilibrium. My starting point here was let's take something that's a dominant-strategy equilibrium.

Maybe there's no dominant strategy equilibrium. That's OK. But I'm saying if there is one, it must also be a Nash equilibrium.

And now let's go see if we can do this next step. So part 2 is to show that every Nash equilibrium is rationalizable. So if I start with something that's Nash equilibrium, it better be a rationalizable strategy profile.

Now this seems a bit hard because the way we defined S infinity was really tricky. We went through this whole iterative procedure. We had an algorithm. We had to go through infinitely many steps.

So do you have any ideas about how we might show this? Any mathematical technique that might be useful for this kind of proof? Yeah.

AUDIENCE: Proof by contradiction.

IAN BALL: We could-- I think a proof by contradiction would work, yeah, but I saw another hand. Yeah.

AUDIENCE: Induction.

IAN BALL: Induction is a good idea. Contradiction would work. A lot of things, you could convert it.

But I think induction is a really natural idea because we've defined S infinity inductively. So we'd like to use induction. And if you haven't heard of induction, that's OK, I'll go through it.

OK, but let's first just understand what we're trying to do. We're going to follow the same procedure as before. So suppose s star-- this is a strategy profile-- is a Nash equilibrium.

What do we want to do? We want to show that s star is rationalizable. So we want to show that s star is in S infinity.

But doing that directly is hard because S infinity is only the result of this iterative procedure. But remember, from over here, the way we got S infinity is we just kept making the set smaller and smaller and smaller. So one way-- or what's equivalent to showing that s star is in S infinity-- so I'll say, or equivalently and easier to show S star is in Sk for all k.

So let's just say this in words. We want to show that s star survives every round of elimination and makes it till the end. Well, the way we show that is we just show it makes it to the first step and the second step and the third step and the fourth step and the fifth step, and it never gets thrown out.

And therefore, it must survive all the way to the end. So we want to show that this strategy profile must survive every round of elimination. And since we're showing something for every k, this is where we're going to use induction.

So what do we want to show? We want to show that, first-- so there's a base case of k equals 0. We want to show first that s star is in there after the zeroth round. And then we have an inductive step that says if it's made it to the k-th round, It's also going to make it to the k plus first round.

And then we can just tie it all together. If we know it makes it to the zeroth round, then it must make it to the first. And if it makes it the first, it must make it to the second and so on. And we can just chain our reasoning together.

So for the base case of k equals 0, what we need to show is that s star is in S0. And maybe I'll put a question mark because this is not a claim. I'm not saying it's true. It's what we're trying to show.

So is s star in S0? Well, yeah, because S0 is just the set of all strategy profiles. And s star is a strategy profile, so that's pretty easy.

This is just by definition. The heart of the proof is the inductive step. So here's the inductive step.

What we want to show is, in general, is if s star makes it through the first k rounds, then it must make it one more round. It must also make it to the k plus first round. So what we Want To Show-- and maybe I'll say WTS because it's always important it proves to be clear about what you're assuming and what you're trying to show. So WTS is just an abbreviation for Want To Show.

What I Want To Show is that if s star is in Sk, then s star is in Sk plus 1. So if I haven't thrown it out by the k-th round, I'm not going to throw it out at the k plus first round.

I think if we just look at it like this, it can be a little tricky because, remember, the way we defined this procedure was we had to reason player by player. So let's write that more explicitly. What does it mean for this strategy profile to be in Sk?

It means the i-th component of that profile, player i's strategy is in si k for each of the players i. And I think that's a bit more concrete. So that means that si star-- that is, remember, s star is a profile of strategies.

si star is the strategy that player i is actually using. That must be in si k. And this is true for all i.

And now, symmetrically, what we're trying to show is that si star is in si k plus 1 for all i. Now, when we go through proofs and arguments, we only go through the right way of doing it, and I think you can sometimes lose sight of how hard it is because you don't show all the wrong paths. And there's an easy mistake people make here.

What's really tempting at this point is to say, OK, let me reason player by player. So I want to look at player i, and I want to say if player i's strategy has survived up to round k, then player i strategies should survive up to round k plus 1. That's not going to get you very far.

And the reason is that which strategies for player i survive up to round k plus 1 depend on which strategies of player i's opponents have survived up to round k because, when I'm choosing, when we go through the algorithm, what's the best response for player i? It depends on what strategies of my opponents have survived so far.

So we really have to reason simultaneously about all the players. We can't just reason through it separately because which strategies survive for player i depend on which strategies of the other players have survived so far. So let's try to go through this argument.

OK, so here's the argument. So let's consider player i.

So we know that si star is in si k. But I've argued that that's actually not so useful for us. What's more important from the perspective of player i?

What do I care about as I'm player i? I care about what strategies my opponents might play. So the important thing is that we know that s negative i star is in s negative i k.

This is the key part of our-- so our inductive assumption is that everyone's strategy has survived so far? But the relevant aspect of that assumption is the fact that all of my opponent strategies have survived so far because that's what's going to determine what strategies are best responses for me. So this is what we know.

And what we want to show is that si star is in si k plus 1. So this is what we know. This is what we want to show. I think this is often a good way to try to get through proofs.

So I think the first thing we should do is just write down some definitions. What is si k plus 1? What is the definition of si k plus 1? Yeah, over here. In the Boston sweatshirt, yeah.

AUDIENCE: It's like all-- it's the set of strategies that are a best response to something in si k.

IAN BALL: Exactly right. And we have a notation for that. That's exactly what it is in words. And in math-- we'll just convert that into math, is s negative i k.

So which strategies survive up to the k plus first round? It's all the strategies for player i that are best responses for player i to some belief over the strategies in s negative i k? So these are the strategies I consider it possible that my opponents might play. And I want to think about all best responses to that.

So now we just have to put things together. We know that s negative i star is in s negative i k. We want to show that si star is a best response to something in s negative i k.

What's the key step? What's the leap here? So why is this true? Yeah.

AUDIENCE: So si star is by definition a best response to s negative i star, and that's what we just showed because their game best [INAUDIBLE].

IAN BALL: Exactly, right. So by the definition of Nash equilibrium-- let me explain this. By the definition of Nash equilibrium, we know that si star is a best response to s negative i star.

And remember, this is a small s, this is a big S, this is a small s. So we know that si star is a best response to s negative i star. And we know by our inductive hypothesis that s star is a strategy in s negative i k.

So it just follows by definition that, therefore, s negative i star is a best response to some belief about things in s negative i k. Specifically, it's a best response to the belief that puts all the probability on this strategy. So we're done.

We've shown that si star is in si k plus 1. And if we reason exactly the same way for all the players, we'll show that this holds for every single player i, and now we've completed our inductive step. So again, I think each individual step is not too hard. But piecing it together and thinking through this is a bit tricky. So again, are there any questions on this argument?

So now let's go a little farther. Let's go to another example.

This is another classic game, and this is called hide-and-seek. So think of this as the childhood game. Someone hides, the other person seeks, and the hider doesn't want to get caught, and the seeker does want to catch the other person.

Of course, there's more substantive examples. You might think of someone cheating on their taxes, trying to hide, and then the IRS seeking, trying to figure out which line of their taxes they cheated on. So I think it has broader implications, but we'll keep with the hide and seek story.

So let's say player 1 is the hider, and there's two places they can hide. And player 2 is the seeker, and there's two places they can look for the person. So the seeker is going to win if they look where the person hid. So that means we're going to have-- give myself a little more space-- negative 1, 1, negative 1, 1.

Because here, person 1 hid in spot A, person 2 looked for them in spot A. That means person 1 gets caught, so they get a payoff of negative 1, and the other player gets a payoff of 1. And the same is true if they both hide and seek in spot B.

But then, over here, we're going to have the opposite. So here, if i, as player 1, hide in spot B and player 2 looks for me in A, then they're not going to find me. And therefore, I get a payoff of 1 and the other player gets a payoff of negative 1. You could also think as the hider as the penalty taker and the seeker as the goalie. There's a lot of different stories you could tell.

So notice a difference between this game and the game we discussed before. So this is a game of pure conflict of interest.

In the Boston game, there were both. There was some conflict of interest because one person liked the Red Sox and one person liked the Celtics. But there was also some common interest because they both liked each other's company, and they wanted to be together.

Here, it's purely a conflict of interest. I want us as player 1. I want us to hide in different places. So maybe I'll say that. So what does player 1 want?

They want them to look and hide in different places, maybe different locations. And player 2 wants them to be in the same location. Of course, the part is you can only choose your location. If you chose the other person's location, then it would be easy.

So as a first step for this, let's go through our standard procedure where we compute player 1's best responses to player 2 and player 2's best responses to player 1. So let's do it slowly. If I'm player 1, if player 2 guesses A, then I want to hide.

So I want to do B. And if player 2 guesses B, I want to hide and be in a different place. So I would like to be A.

Now let's look at it the other way. Now let's look at the perspective of player 2. If player 1 plays A, then player 2 does better if they also go A because they want to find the person. And if player 1 goes B, then player two wants to go B.

So we've computed the best responses. We can look, and if we follow our algorithm, where are the Nash equilibria in this game? We want to look for a cell where both numbers are underlined.

But here, we have a problem because none of the cells are underlined. So what this tells us is that this game has no-- and this will be important in the second-- pure-strategy Nash equilibrium. And we can think through it reasonably right.

In order for us to have a Nash equilibrium, neither player can have a profitable deviation. But if we're in the same place, then someone certainly has a profitable deviation. The hider would rather deviate if we're in the same place.

And if we're in different places, then, again, someone would like to deviate, the seeker would like to deviate and change places so that they find the person who's hiding. So because of this immense conflict of interest, we can't have a pure-strategy Nash equilibrium because someone's losing, and whoever's losing would rather change their behavior so that they could win. This is a zero-sum game in a formal sense that we'll talk about more later.

But we'd still like to make a prediction about play in this game. If we just use rationalizability, well, anything could happen because I can hide an A if I think the other person's going to B, and I can hide in B If I think they're going to A, and we can explain everything. So I don't know, how would you play this game? If you were taking penalty kicks in the World Cup, how would you-- what strategy would you use? Yeah.

AUDIENCE: Mixed-strategy.

IAN BALL: You'd use a mixed strategy, right? So you'd randomize. And in this case, how do you think you would randomize? Would you play A 99% of the time?

AUDIENCE: 1/2 and 1/2.

IAN BALL: Probably 1/2 and 1/2 because you don't want yourself to be predictable. If you look at the data on World Cup penalty takers, that's a pretty good prediction. I mean, sometimes right-footed players have a bit of a left advantage, and there's all these subtleties.

But roughly, people mix 50/50 because if I develop a reputation, if everyone knows that-- if you think about playing hide and seek with a young kid, sometimes they just always hide in the same place. It's pretty easy to beat them. And even if you don't always hide in the same place, if you hid in the same place 95% of the time, you'd also be pretty easy to find. So it seems pretty intuitive that you'd want to do 1/2, 1/2.

So that's a good guess. And that's right. So it turns out it will have a mixed-strategy Nash equilibrium. But 1/2, 1/2, meaning player 1 randomizes 1/2, 1/2 between A and B, let me write this a bit more explicitly.

1/2 A plus 1/2 B comma 1/2 A plus 1/2 B. This means player 1 mixes 50/50 and player 2 mixes 50/50. This is going to be a Nash equilibrium in mixed strategies.

So I'd say this is-- maybe I'll use the acronym Mixed-Strategy. NE, MSNE, mixed-strategy Nash equilibrium. But I haven't even defined what that is. So let's now define that. I think it's exactly what you'd think.

So definition, a mixed-strategy profile-- and now we'll use sigma instead of s to remind us that this is mixed. So sigma star equals sigma 1 star, sigma n star is a Nash equilibrium-- and we'll give exactly the same definition, for each player i, what happens?

Well, I'm comparing now what I'm supposed to do as player i to a deviation, which I'll maybe call sigma i prime. So this is the mixed-strategy I'm actually playing. This is the mixed-strategy I could deviate to. And I want this to be optimal given the way my opponents are playing.

So this is some new notation. We've been using capital SI to denote all the pure strategies of player i. It's often convenient to use, this, it's the summation symbol, and this is capital Sigma, and this is the set of all mixed strategies by player i.

So we use lowercase s for pure strategy, uppercase S for a set of pure strategies, lowercase sigma for a mixed strategy, and uppercase Sigma for a set of mixed strategies. So it's the same idea, given the way my opponents are playing, my strategy is a best response. My mixed strategy does weakly better than any other mixed strategy I could use.

But if you look at this definition, it's pretty messy because there's so many different mixed strategies I could deviate to. Even in a game with just two strategies, I could deviate to any probability between 0 and 1. That's a lot of deviations to deal with.

But fortunately, we don't actually have to worry about that. So it turns out, even if I were to check every mixed strategy, it's actually enough just to check my pure strategy deviations. So it turns out this is the same as-- the same thing, but with si prime for all si prime. So when I use the two dashes, I'm just saying everything else in the definition is the same. But instead of saying, it must do better than any mixed strategy, I'm saying it must do weakly better than any pure strategy.

And we could give a formal math proof, but let's just think of the intuition. If I'd rather do what I'm doing than deviate to any pure strategy, then by choosing to flip a coin and mix over those pure strategies, I can't do strictly better because if I mix over those pure strategies, my payoff is just going to be an average of my payoff from all the pure deviations. So if each of the pure deviations makes me weakly worse off, then a mixture of them must also make me weakly worse off.

That's the intuition. You could write down a formal mathematical proof. But the point is, it's really important here that what I am doing can be mixed. So I want people to really understand this. This is crucial.

It's crucial that we're allowing my equilibrium strategy to be mixed and not pure. But when I contemplate the deviations, it's enough to only look at pure strategies. And that's a key difference here.

So indeed, how can we tell that mixtures matter? Well, already in this game here, we saw that it didn't have a pure-strategy Nash equilibrium. But it does have a mixed-strategy Nash equilibrium. So the fact that we allow the equilibrium strategy to be mixed is essential. But when you're contemplating deviations, you only have to worry about pure strategies.

And let's take that randomization idea little further. Why would I ever want to play a mixed strategy? I mean, why am I willing to flip a coin and sometimes go to A and sometimes go to B? I mean, if A is better, wouldn't I just always want to go to A?

And if B is better, wouldn't I just always want to go to B? Why would I ever want to flip a coin? Or why would I ever be willing to flip a coin and play A some of the time and B some of the time? What must be true? If I'm-- yeah.

AUDIENCE: A is better in some circumstances while B is better [INAUDIBLE].

IAN BALL: OK, that's true, but what about given my belief, given the way my opponents are playing, what must be true? If I strictly preferred A to B, given the way my opponents were playing, would I want to mix between A and B? No. Yeah.

AUDIENCE: Does your opponent also have to be playing a mixed strategy?

IAN BALL: So that's related, but maybe not quite where we are. So if I'm willing to mix between A and B, it must be that I'm indifferent between A and B. That's the key thing.

If I strictly preferred A to B, I would never play B with positive probability at all. And if I strictly preferred B, I would never play A with positive probability? So I'm only willing to flip a coin and sometimes play A and sometimes play B if I'm exactly indifferent between A and B.

Now this gets to your point. Why will I be indifferent between A and B? It'll be precisely because my opponent is mixing. So this is a key observation.

Maybe I'll start a new board.

So this is the key observation. So mixing is optimal. Only if you're indifferent between the things you're mixing over.

And in fact, we can say something even stronger. It's not just enough that I'm indifferent, but I'm only willing to mix over pure strategies that are best responses. So it must be-- and in fact, it's only optimal to mix over pure best responses.

So what does that mean? Well, it's just a silly example. What if I had A, B, and C and given the way my opponents were playing, A gave me 1, B gave me 1, and C gave me 2.

Do I want to mix over A and B? Well, no, even though I'm indifferent between A and B, that's necessary, but it's not sufficient. I certainly don't want to mix over A and B because C is strictly better.

So if I'm willing to mix, I must be indifferent. But sometimes people get confused, and they get it backwards and they say, if I'm indifferent, I must be willing to mix. But no, it's not just enough to be indifferent. It also has to be optimal what you're mixing over.

And that's precisely what I'm saying here. It's only optimal to mix over pure best responses. So if I added D at 2, now, indeed, I'd be willing to mix between C and D.

So now let's try to check-- going back to this game-- that 1/2, 1/2 is actually a Nash equilibrium. So I think the best way to do it is to say, well, let's take the perspective of player 2. And let's suppose that player 1 mixes 1/2, 1/2.

What payoff does player 2 get from playing A and playing B? So what I'm going to imagine is I'm going to add a new strategy here, 1/2 A, 1/2 B. So I'm adding a new row to this matrix that's going to represent the payoffs if player 1 chooses this mixed strategy. So in here, what I want to write is, what are the expected payoffs if player 1 uses this mixed strategy and player 2 plays B? Sorry, plays A.

And then, over here, what am I going to write? I'm going to write the payoffs-- if player 1 plays this mixed strategy and player 2 plays B. And it turns out, this-- well, let's go through it.

If player 2 plays and player 1 plays A with probability 1/2 and B with probability 1/2, then what is the payoff of-- what are the payoffs going to be for the players? Well, player 1 is going to get negative 1 with probability 1/2 and 1 with probability 1/2, which just gives them 0. And player 2 is also going to get 1 with probability 1/2 and negative 1 with probability 1/2. So we just get 0, and if we go to this side, we're actually just going to get 0's again.

And mathematically, the way we form this row of the matrix is we just take this row, multiply it by half, take this row, multiply it by half, and add them together. And when we do that, we exactly get 0, 0, 0, 0. So in particular, the key thing about this is that these numbers are the same.

So what this tells me is if player 1 plays 1/2 A plus 1/2 B, then player 2 gets a payoff of 0 if they play A and a payoff of 0 if they play B. And that means that player 2 is indifferent between A and B, and in fact, A and B are both pure best responses. So if I'm indifferent between A and B, then I'm also willing to mix over A and B. And therefore, it's a best response for player 2 to mix 1/2, 1/2.

So technically, what I've shown is if player 1 is mixing 1/2, 1/2, it's a best response for player 2 to mix 1/2, 1/2. So technically, if I want to finish the argument, I need to show the other direction, that if player 2 mixes 1/2, 1/2, then it's a best response for player 1 to mix 1/2, 1/2, but it's going to be entirely symmetric. I could add another column here and go 0, 0, but we can check that-- I'll just put a check mark.

We checked it. It's an equilibrium. Don't do that on your homework, but that's all we need.

Now, let's return to the Boston game. I don't know if we have it. I may have to write it again. You can write it again. That's OK.

So now let's go to the BOS game. This is a bit trickier. And we have Celtics, Red Sox, Celtics, Red Sox 2, 1, 1, 2, 0, 0, 0, 0.

So this game, we already found some pure equilibria. We found two of them. We said there's one equilibrium where we both play C, and there's one equilibrium where we both play R.

But it turns out, there's also a mixed-strategy equilibrium. So here, what we'd like to do is find a mixed-strategy Nash equilibrium of this game. Now, of course, we could say the pure-strategy Nash equilibria we already found are technically mixed-strategy Nash equilibria where people's mixing probabilities are just 1 and 0. But what I mean is something that's actually mixed.

So what I really mean is a mixed-strategy Nash equilibrium that's not a pure-strategy Nash equilibrium. OK, well, I think the first step is let's just write some notations. So we're going to look for an equilibrium. We're going to look for some mixing.

So let's look for an equilibrium of the form-- well, player 1 is going to put, say, probability P1 on C. And let's say player 2 is going to put probability P2 on C. So our goal is to find a Nash equilibrium of the form player 1 plays P1 times C plus 1 minus P1 times R. So what this means is, with probability P1, player 1 goes to the Celtics game, and with probability 1 minus P1, player 1 goes to the Red Sox game. And then, similarly for player 2, but we have potentially a different number here.

And we want to find an equilibrium of this form where maybe what we'd like is we want it to actually involve mixing. So let's say P1-- 0 is less than P1 is less than 1, and 0 is less than P2 is less than 1. This is what we want. We want something where each player is actually mixing. They're not just playing one strategy with probability 1.

So I guess the trick is to think, well, what must be true here? So if we want this to be a mixed-strategy Nash equilibrium, it must be that player 1 is willing to mix between C and R, given the way that player 2 is playing. And based on what we said before, what does that mean? What must be true if player 1 is willing to mix between C and R? Yeah.

AUDIENCE: Player 1 [INAUDIBLE] C, R [INAUDIBLE].

IAN BALL: Exactly, and then the other direction. Player 2 must also be indifferent between C and R, given player 1 strategy. So the way that we as the analyst often solve for an equilibrium like this is we look for mixing probabilities that make the players indifferent. So let's look at this, there's two steps.

So the mixture P1 must make player 2 indifferent.

And this is a key point. You might think I made a typo here, but I didn't actually. It turns out that the probability with which player 1 mixes is what determines whether player 2 is indifferent because player 1's mixing probabilities in equilibrium determine the beliefs that player 2 has.

So whether player 2 wants to play C or R depends on their beliefs about player 1 and the frequency with which player 1 plays C or R, so it's, actually, player 1's mixing probability that is responsible for determining whether player 2 is indifferent. And conversely, it's player 2's mixing probability that determines whether player 1 is indifferent.

So when you try to solve for these equilibrium, it all looks nice. We have two indifference conditions and two numbers to solve for. Mathematically, it's pretty easy, but the thing to remember is that player 2's indifference condition is actually, in this case, going to be determined by P1, and player 1's going to be determined by P2.

So player 2's indifferent condition, we need to make sure that player 2 gets the same payoff from C as from R. So let's look at player 2. If player 2 plays C, well, they get a payoff of 1 if player 1 plays C and a payoff of 0 if player 1 plays R.

So what is their expected payoff? Their expected payoff from playing C is going to be P1 times 1 plus 1 minus P1 times 0. But the 0 disappears, and we just get P1 in the left-hand side. So here, P1 is the utility, the expected utility for player 2 from playing C.

So maybe I'll put a C here and an R here just to keep track of what we're comparing. And this makes sense. If I'm player 2, and I go to the Celtics game, I know that if my opponent goes to the Red Sox game, I get a payoff of 0.

So my only utility comes from if they also go to the Celtics game, that happens with probability P1. And if it does happen, I get a payoff of 1 because I like the Red Sox more-- I only get a payoff of one from the Celtics. What about the other case? If we have the R-- if I play r, what is my payoff going to be? Yeah.

AUDIENCE: [INAUDIBLE] 2 times 1 minus R.

IAN BALL: Exactly, right. So here, again, I get no payoff if they go the other way. 1 minus P1 is the probability that my friend also goes to the Red Sox game. But now there's a 2 here because I really like the Red Sox. Or if they go to the Red Sox game with me, I get a payoff of 2 rather than 1.

And now let's do the same thing for player 1. I'm player 1. If I play C, given that my opponent is playing C with probability P2, what is my payoff going to be as player 1?

Well, it's going to be 2. I get 2 with probability P2 and 0 with probability 1 minus P2, so I just get 2 P2. And then if I play R, again, 0 with probability P2, 1 with probability 1 minus P2, so I get 1 minus P2.

And we can just do some algebra here. And if we pull this over, we're going to get P2 equals 1/3. And then P1 equals 2/3.

So I want to caution here against, I think, a common misinterpretation of this. What a lot of people say is they say, OK, this makes sense. Player 1 goes to the Celtics game more often.

P1 is 2/3, more than 1/2, because player 1 likes the Celtics. And player 2 goes to the Red Sox game more. Player 2 goes to Celtics game only with probability 1/3, which means they go to the Red Sox game with probability 2/3 because player 2 likes the Red Sox more.

Is that right? Do you see a problem with that reasoning? I don't know, what do you think? Does that seem right to you?

So it turns out that the reason, if we really look through the equilibrium, the reason that player 1 goes to the Celtics a lot is not because they like the Celtics more but because their opponent likes the Celtics less. Player 1's mixing probabilities have to make player 2 indifferent. And because player 2 doesn't like the Celtics very much, they're only willing to go to the Celtics game if they believe that player 1 is very often going to the Celtics game.

And let's do the reverse reasoning. Why does player two go to the Red Sox game a lot? It's true that player 2 likes the Red Sox a lot, but the real reason for this equilibrium is that they have to go to the Red Sox game a lot to make player 1 willing to go to the Red Sox game because player 1 prefers the Celtics, so they're only going to go to the Red Sox game if they think it's more likely than not that player 2 goes to the Red Sox game. So I think that's just a key point here about what's driving the structure of this equilibrium.

And I think that should make you wonder a bit because I shouldn't care about making my opponent indifferent. So I think an issue, maybe a question, is if I'm player 1 here, given the way my opponent is playing, I am indifferent between C and R. So why should I mix with exactly the right probability that makes my opponent indifferent?

It doesn't really seem like, if I'm indifferent, I mean, why go through the difficulty of doing all the randomization? And I think that's a question we'll just leave for now. We'll talk a bit more throughout the course about this, but I think it's something to keep in mind, that it's true, I'm willing to mix, but why do I actually mix?

We do see it in practice. I don't know. I will open it up. Are there any thoughts?

I mean, why do you think poker players actually do mix? They're indifferent, they should be indifferent if everyone else is mixing just right. Why do they mix?

Why would you mix if you took penalties? What would go wrong if you didn't mix? So yeah.

AUDIENCE: They would know what you were going to do, and they would update their strategy based on that.

IAN BALL: Right, so often, the story we tell about mixed strategies is often really a story about a dynamic process where the game is repeated. And for this reason, we might think that a mixed-strategy Nash equilibrium is a more compelling prediction if the game is repeated over time. You don't want to have a reputation as someone who always shoots left because then your opponents would update on that.

Now, then we have to really think about, well, what does it mean to develop a reputation for something? And that takes us down a whole new path that people have tried to formalize and think through. It's not clear exactly what it means. So that's one story.

I guess, another story I would say is sometimes mixing is not literal mixing, but what mixing means is you're making a decision based on something that the other person isn't aware of. So it may be that-- let's take, I don't know, how someone dresses. You may think, oh, if someone just randomizes the way they dress, but really, they wear certain colors on certain weather because that makes them feel happy.

But to you, not knowing their careful algorithm for how they choose their clothes as a function of the weather, it just appears random to you. Or similarly, maybe a penalty kicker, when they come to the spot, feels a direction. One of their feet feels a little different, they just feel a direction to go to, and then they make a deterministic choice about which direction to shoot, but because you, as the opponent, can't see that signal that they're basing their shot on, it appears random to you. So sometimes we literally observe randomization, and sometimes behaviors appear random because they're based upon things that we can't actually observe. So just one thing to keep in mind.

And I think that's good. Let me finish with one last-- two last things. So notice that, in the game that we looked at here, this game did not have a pure-strategy Nash equilibrium.

But when we moved to mixed-strategy Nash equilibrium, it did have a mixed strategy Nash equilibrium. And the question is, could we come up with some really weird game where there's no equilibrium. That maybe wouldn't be so very good for our theory because then we wouldn't have a prediction about how people play.

And fortunately, it turns out-- and this is really Nash's key result in 1948-- that that's never going to happen. Every game, within reason, has a mixed strategy Nash equilibrium. So more precisely, let me say every finite strategic-form game-- would be a little more precise, I'll say-- has at least one.

So at least one. It might have only one. It might have two. It might have three. It may have infinitely many, but it has at least one.

And when I say mixed strategy, that includes pure strategy as a special case. So I might say mixed or pure, but remember, we're going to think of a pure-strategy Nash equilibrium as a special case of a mixed-strategy Nash equilibrium where the players mix with degenerate probabilities. Now, this is not something we're going to prove in the course.

This is actually quite a deep result. So remember, a couple of weeks ago, last week, we talked about this dominance/best response duality, and I also didn't prove that. That is somewhat deep.

This is much deeper. So this is maybe the mathematically deepest result in the course. And it uses something called a fixed-point theorem.

And if you want to google it, it's either sometimes attributed to Brouwer or Kakutani. And these are, actually, quite deep mathematical results. So one of the mathematical results was actually proven, I think, in 1944.

And then Nash proved this result in 1948. So it was really using cutting-edge math to prove this new result. Maybe not so cutting edge anymore in 1944, but at the time, the gap between the math frontier and the econ frontier was actually very, very close.

If you want to see the proof, you can look in the lecture notes. The lecture notes give a proof building upon these theorems. The real depth is proving these theorems, which we're not going to talk about. We're not going to talk about this proof either, but just want you to keep that in mind.

And before we go, we have three more minutes. So let me just conclude with a few interpretations of Nash equilibrium. I think we talk about Nash equilibrium as the crowning achievement of game theory. But over the last seven years, people have always been searching for justifications of Nash equilibrium.

And I would actually argue that Nash equilibrium does not have as compelling of a justification as, say, rationalizability does. Of course, it makes stronger assumptions, and it gets sharper predictions, but whether we actually find it compelling, I think, is not totally obvious. Again, on exams, if I ask you to solve for a Nash equilibrium, you have to solve it. You can't just say, I don't believe in Nash equilibrium, but I do think it's helpful to think about how compelling this is.

So I want to present a few interpretations of this. So one interpretation is about a stable convention. So the idea of Nash equilibrium is really about stability.

It's that given the way people are behaving, no one wants to deviate. So a great, I think, example of this is the side of the road that we drive on. Why do we all drive on the right in the US and drive on the left in England? You might say, oh, you get a ticket if you drove on the left, but that's not really the reason. It's really because you're afraid you'd get hit by a car.

So if everyone drives on the right, it's best for you to drive on the right. If everyone drives on the left, it's best for you to drive on the left. These are both stable conventions, and we don't really need to ticket people for driving on the wrong side of the road. I mean, we do, but that's not really what motivates them, I think, to drive on the correct side of the road. Another is you can think of it as a self-enforcing agreement.

Again, self-enforcing because no one wants to unilaterally deviate from the agreement. And I think this explains the success of a lot of multilateral international agreements. When countries sign treaties, there's no court that can adjudicate a treaty between countries, or generally, there's not.

So why do people follow the treaty? Well, often when treaties are set up in the right way, people don't have an incentive to unilaterally deviate. And even though there's not some court that would throw a country in prison, they often are still motivated to follow these agreements.

And then another interpretation is as a steady state. And this gets to a discussion of penalty kicks. If you were to always shoot right, the goalies would start always guessing that you were going to shoot right. So then you might start shooting left, and then, over time, we think that you might converge to the steady state.

Whether there's actually a real justification for dynamics converging is actually a bit more subtle. But the point is, if you are in a Nash equilibrium, it's steady, and there's no reason to move away from that. And that's the final interpretation, and I'll leave it there. So I'll see everyone next week.