[SQUEAKING] [RUSTLING] [CLICKING]

IAN BALL: So today we're going to continue with what we did last week, and we're going to talk about something called rationalizability. And I'll say up front that I think today is the first class where maybe there'll be a bit more notation, a bit more math. So if anything's unclear, if some notation isn't clear, feel free to ask me questions because there's probably other students who also find it unclear.

So first, let's remember or recall how we defined a strategic form game. So remember, a strategic form game specified the strategy sets for all of the players, and also the payoffs for all of the players. So one way I'll often write it is u from S1 Sn to Rn. So what does this mean?

This u function assigns to every profile of strategies some vector of payoffs that specifies each of the n player's payoffs. So with this notation, we might think of a strategy profile S is going to get mapped to a vector u of s, which is just all the payoffs the players get put together in a vector.

So we have u1 of s, all the way up to u of s. This tells us player one's payoff, if the profile of strategies that's played is s, all the way up to player N's payoff, if the profile of strategies that's played is s.

And I'll point out here that we're generally going to use lowercase u throughout the course. So in the first couple of lectures we made this big deal, which I think correctly so, about the distinction between the von Neumann-Morgenstern utility, which was defined over consequences, and then expected utility. That was capital U. And that was defined over lotteries, over consequences.

But throughout the course, as we start talking about beliefs and mixed strategies, it gets a little cumbersome to think every time, should this be capital U or lowercase u? So I'm throughout the course from here on out going to always use lowercase u, with the understanding that that can represent expected utility if there's any beliefs or mixed strategies or moves by nature in the extensive form. But as we get to more complicated things, it's just a little less clear when it's big U and little u. So I'm just going to stick to little u as long as we understand the conceptual distinction.

So what we talked about last time is which strategies were reasonable to play? So last time, we talked about strategies that were reasonable, and I put reasonable and unreasonable in quotes because I don't mean these in really formal terms. I'm just being a little informal here. But let's remember we call this strategy reasonable for a player if it was a best response to some belief. So I'll say BR for Best Response to some belief.

And let's remember what a belief is. A belief that player i has is a belief about the strategies chosen by all of the other players. So whenever I talk about player i's belief, that's a belief about the partial profile of strategies played by the other players.

And then we alternatively said we can think of a strategy as unreasonable if it is strictly dominated, which we defined last week. And crucially, it wasn't enough to just think about strict dominance by another pure strategy. We had to consider strict dominance by mixed strategies, so if it's strictly dominated by some mixed strategy.

And what we showed last week is that these are two sides of the same coin. Either a strategy is a best response to some belief, or it's strictly dominated by some mixed strategy. But every strategy falls into exactly one of these two categories.

So you can take a positive view and say, I only want to play a strategy that's a best response to some belief. Or you can take the negative view and say, I don't want to play a strategy that's strictly dominated by some mixed strategy. And either approach you take is going to give you the same set of strategies.

So today, we want to be maybe a little more formal. I don't really love this terminology, but we'll use it, and we'll say what it means for a player to be rational in this context. We'll say player i is rational-- well, if player i plays a strategy that's in this good collection of strategies, or equivalently, is not in this bad collection of strategies.

So I'll say player i is rational if she plays a best response to some belief. That is, the strategy she chooses is a best response to some belief about the other player's strategies.

I want to be a little careful using the term "rational." This is a word that's used in a lot of different ways in a lot of different contexts. But in this specific context of game theory, this is usually the term that's used. So I'm going to use it. But I'll warn you that in other game theory courses you might take or other things you see, rationality can be defined to mean different things because it's an important, difficult concept.

But let's look at this definition. And it's restrictive in that the player must be playing a best response to some belief. So the best response, part of it is what imposes some restriction. But notice, there's no restriction on the belief here.

So the assumption or the restriction of rationality by itself is a very weak requirement because it doesn't constrain or discipline the beliefs that player I forms about the other player's strategies. So this is purely about player I's rationality, and it says nothing about the other player's rationality.

But let's think a little bit and let's say what if not only is player I rational, but what if player i knows or believes that the other players are rational too? Would this change the way that player i might play, or would this impose any additional restrictions? What do you think? What might change if player i, in addition to being rational herself, also knows that the other players she's playing against are rational? Yeah.

STUDENT: If you change your belief, that now she knows what the other players [INAUDIBLE] in this [INAUDIBLE] strategy [INAUDIBLE].

IAN BALL: Exactly. Right. So-- do you want to say more or-- yeah. Yeah. Yeah, great. So I just didn't want to cut you off. Exactly. So before we had no restrictions on beliefs. We said, all player i is going to do is they're going to play some best response to a belief, but we don't know where that belief comes from.

But now if player I knows that the other players are rational, they shouldn't form a belief that assigns positive probability to a really weird, bad strategy that a rational player wouldn't play. So the idea here is that this is going to restrict or constrain the beliefs that player i should or reasonably would form about others-- constrain beliefs about other strategies. And do we stop there? Are we done? Or can we keep going? Yeah?

STUDENT: I guess the other players from their perspective would update based on [INAUDIBLE] player I's [INAUDIBLE].

IAN BALL: Right. So you might think, OK, not only is it relevant if I know that the other players are rational. But do I know that the other players know that I'm rational? Because if the other players know that I'm rational, that's going to affect the beliefs that they form about my strategy, and therefore, further restrict the way they might play, and therefore further restrict the beliefs that I would form about them.

Now, I think you can see pretty soon this is going to get pretty complicated, and just writing it in words is going to become cumbersome. So now we're going to try to take a formal approach to this kind of reasoning that captures what's often called higher order knowledge of rationality.

So first order knowledge of rationality means I know other people are rational. Second order knowledge is I know that they know that I'm rational. And third order I know that they know that I know, and we can see how this goes on and on and on. And we want to capture these higher degrees of knowledge about rationality in order to make more precise or sharper predictions about playing the game.

So let's go through-- this is a bit abstract. So let's go through an example to try to convey this idea. And this is going to be the same example we did last class, but we're going to add in the payoffs for player two as well.

So we're going to have a game like this. Remember player one is always the player who chooses rows, and player two is always the player who chooses columns. We always put player one's payoffs first. So we're going to have 2, 0, negative 1, 1. Let me copy from my notes here.

So last week, we used these same payoffs for player one, but we didn't specify any payoffs for player two when we talked about this game. And just as a reminder, we represent the game in this nice normal form, what's sometimes called a bimatrix, because it's a matrix of numbers, but each entry has two numbers, so "bit" for two.

But let's remember what's going on here. S1, the set of strategies for player one, are just the rows of this bimatrix. So we have top, middle, and bottom. And S2 for player two is just left and right.

So let's try to start going through this procedure of rationality. So let's first think about, well, if player one is rational, what strategies might they play? And a good way of going about this is sometimes helpful is to say, well, we know that rationality is a bit complicated because we have to consider all possible beliefs about the other players. But let's start with the simpler beliefs, where we know what strategy the other player is playing.

So if I'm player one, and I know that my opponent is playing L, then I can just look down and say, well, I get 2 if I play T. I get 0 if I play M, and I get negative 1 if I play B. So the best response for me as player one is to play T. And we often denote that by underlying the first entry here.

So the fact that we're underlying player one's payoff means we're thinking about player one's best response to player two's strategy. And the fact that it's in the first column tells us we're thinking about player one's best response to player two playing L.

Now let's do the same thing if player two plays R. If player two were to play R, then player one knows that they get negative 1 if they play T, 0 if they play M, and 2 if they play B. 2 is the highest number there, so we're going to underline this.

And we can immediately see two strategies that are reasonable for player one-- T and B. Now, the trickier part is, are there any more? Does anyone remember?

So M here is not a best response to L or R, but we haven't actually checked if it could be a best response to a belief where I think you play L with some probability and R with some probability. Does anyone remember? Is M best response to any belief that player one has? You're shaking your head. The answer is no. And why? How do we know that? Yeah?

STUDENT: Because of the theorem that says either a strategy is a best response to some belief or is strictly dominated.

IAN BALL: Exactly right. And you're saying that this is strictly dominated by something? It is. It is going to be strictly dominated by something. This is what we went over last week that was a bit of a while ago.

We said that M is strictly dominated by the mixture of playing T with probability 1/2 and B with probability 1/2. That's something we checked last week. So we know that the only, quote unquote, reasonable strategies for player one are T and B.

So let's maybe think of that step one. And when I say step one, I'm going to say, let's assume or suppose that both players are rational. And just from that assumption alone, what can we conclude about the strategies that might be played? So P1 and P2 are both rational.

Well, if they're both rational, we've said that that tells us something about what player one can play. So I'll write S1 prime here to say, let me collect together the strategies that player one could play, or might reasonably play, if player one is rational. And in this case, we just get T and B.

And now let's try to do the same for player two. So let's go through the same procedure. Suppose I'm player two. Imagine player one is playing T. Then I get 0 if I play left, and 1 if I play right. Right is better. So I'm going to underline it.

And notice here I underlined the second number because here I'm keeping track of the best response by player two to the beliefs they have about player one. And I'll do the same thing here. If player one plays M, I can get 10 or 0. 10 is higher, so I'll choose L. And if they play B, I can get negative 6 or 0. 0 is higher, so I come here.

And we can immediately see that R is a best response to some strategy, and L is a best response to another strategy. So in this case, we're already done because we already see that both of these strategies are reasonable. Each of them is a best response. In fact, in this case, not only to some belief, but to a belief that puts all the probability on a single strategy. So here, we'll say S2 prime equals L and R.

And notice an important point here. Rationality doesn't always restrict the strategies that a player might play. In player one's case, player one's rationality does throw out one strategy. It threw out M because M was strictly dominated by another mixed strategy and was never a best response to any belief.

But for player two, even if they're rational, we can't say much more. We can't say anything else about what they might play because what's best for them depends upon what player one does. And if player one plays T, it's great for them to play R. And if player one plays M, it's great for them to play L So we can't rule out any strategies yet. So that's kind of step one.

Now let's go to step two and impose a bit more-- slightly strengthen our assumptions. Let's now say in step two, we'll maintain our assumption that P1 and P2 are both rational. But we'll add the further assumption that each player knows that the other player is rational.

So I'll say "and" to remind you that I'm maintaining the assumption up here, each knows that the other is rational. So our numbers are kind of off here. We could call that step zero, but this is first order knowledge of rationality here in step two.

So now let's maybe write S1 prime, double prime, to be the strategies that player one could play. If player one is rational and knows that the other player is rational. So if player one knows that player two is rational, what can they conclude from that?

Well, in this case, actually not much. They know that instead of the player playing something in here, they're going to play something in here. Well, those are the same. So actually in this case, player one's knowledge of player two's rationality doesn't change things.

And we're actually going to get the same answer down here. We're going to get T and B. Because in this case, my knowledge of player two's rationality doesn't actually restrict the beliefs that I form.

But let's do the second player. I think this will be a bit clearer. What about player two? So now that player two knows that player one is rational, how does that change player two's beliefs? Yeah?

STUDENT: [INAUDIBLE] because they know that player one won't play M, so it's so not rational.

IAN BALL: Great. So let's go through that slowly. So if I'm player two, if I know that player one is rational, I can reason through the way player one would behave. And I can say if player one is rational, they're only going to play T or B. Those are the only two strategies they could play.

And then I'm going to look up and I'm going to say, OK, if player one plays T, then I'm better off playing R. And if player one plays B, then I'm also better off playing R.

Now we have to be a little careful here because I might also think, well, maybe they'll play T with some probability and B with some probability. But it turns out even if they played those with some probability, you'd still be better off playing R.

And the reason is that, well, R is strictly better than L given T, and strictly better than L given B. So it's not too hard to check that if they formed a belief about those two things, R would still be better.

So now we go down to R. So what changed when we moved from step one to step two? Notice the set of strategies that we think player two might play, or could play, has gotten smaller, not because player two is rational.

We already assumed player two is rational. But what further restricted our prediction, or sharpened our prediction, was our assumption that player two not only was rational herself, but knew that player one was also rational. And by thinking through how her opponent being rational might play, player two was able to restrict things. So now let's see if we can go one step further.

So now it gets a bit cumbersome to write, but maybe I'll say step three is what? It's that, and each player knows that each other player knows that the other player is rational. This is a bit trickier to say, but-- and knows that the other knows that they are rational. It's a bit sloppy, but--

So if we read through this from player one's perspective, player one knows that the other player, player two knows that they, player one, are rational. And similarly, player two knows that player one knows that player two is rational.

So we might-- if we want to write this in words, it's like one knows that two knows that one is rational, and two knows that one knows that two is rational. It's two different statements here, and the arrows indicate knowledge. This is that two knows that one is rational, and this is that one knows that two knows that one is rational, and so on. And maybe now we'll call this S1 triple prime and S2 triple prime.

Well, S2 triple prime is pretty easy because player two, we've already narrowed it down to the only strategy they have left. So that's just still going to be R. That's not going to change. But what about player one? What is player one-- what is reasonable for player one to play given the second order knowledge of rationality? Thoughts? Yeah? B, right?

Because now if player one knows that player two must play R, then they look up, and they say, well, if player two plays R, I've already actually underlined it. B is my best response to R. So if I know that player two is playing R, then my best response is B.

And notice here the structure of this when I go through the steps. We write S1 triple prime and S2 triple prime next to each other. But what defines S1 triple prime is actually my knowledge that player two is playing a strategy in S2 double prime. And similarly, S2 triple prime is driven by player two's beliefs about what player one is going to play from S1 double prime.

So we're always kind of looking back to the step before. The step before pinned down how my opponent would play. And now adding one other layer of knowledge tells me that I know my opponent will play that way, and that may allow me to further restrict the strategy that I use, this crisscross pattern as we go down. And now we're done.

In this case, there's no strategies left to deal with, and we can check. Not a coincidence here that we have one strategy left at the end, one strategy profile left at the end, BR. And notice that both be that both numbers in that cell are underlined, which reminds us that B is a best response to R, and R is a best response to B, which makes sense and will hold a bit more generally.

So notice here that here, this approach, or we might say rationalizability-- or we haven't defined that yet. But let's say here, we get a unique prediction. When we apply this algorithm, eventually we got to a point where each player only had one strategy left, but that doesn't always hold. So let's emphasize this is not always the case, as we'll see later.

But when this is true, we have a name for it. So when we have a game where we can go through this procedure, and we get a unique prediction for each player's strategy, we call that game dominance solvable. So when this does happen, it's not guaranteed. But when it does happen, we say the game is dominance solvable. Underline this.

Why do we call it dominance solvable? Well, simply by reasoning about dominance-- remember, we often reasoned about best responses, but that's equivalent to reasoning about dominance by our theorem from last week. So all we really needed to use is that players don't play strictly dominated strategies.

They know their opponents don't play strictly dominated strategies. They know their opponents know that they don't play strictly dominated strategies. They know their opponents know that they know that their opponents know that they don't play strictly dominated strategies and so on. And simply by reasoning about dominance, we were able to, quote unquote, "solve the game," in the sense of identifying a unique prediction in the game. Any questions on this so far?

So now we'd like to formalize this. We went through it in an example. But let's introduce the formal algorithm that is going to go through this procedure. So this algorithm is called rationalizability. Or it has two names. So it's either called rationalizability, and rationalize ability is really a statement about the strategies that survive this algorithm.

So in this example, we would say strategy B is rationalizable for player B. And strategy R is rationalizable for player two. Sorry-- player one and player two. But the procedure we use is often goes by a different name, which is the iterated elimination of strictly dominated strategies. And sometimes this goes by the acronym iterated elimination of strictly dominated strategies because it's a bit of a mouthful.

What's the idea? At each step, we implicitly threw out the strategies that were strictly dominated. We eliminated the strategies that were strictly dominated. And we did this recursively, or iteratively. And that's why we call it iterated elimination of strictly dominated strategies.

To formally define this procedure, we need some notation. It was easy in this example to work through it. But to do it more generally we're going to need some general notation. So we really need two pieces of notation to start.

So let's consider player i. We're given player i and some collection of strategy profiles of the other players. So we're given player i, and I'll say S negative i prime, which is a subset of S negative i.

So remember, S negative i is the collection of all profiles, or we might say partial profiles, of strategies of my opponents. And S negative i prime is just an arbitrary subset of those strategies, of those strategy profiles. So it's just some collection of strategies that my opponents might play.

And now we want to define two things. The first thing is BRi of S negative i prime. So what this is going to be is the set of strategies for player i that are best responses to some belief about my opponent's strategies that only puts positive probability on strategies in here. So let's go through this.

Remember, we said we have to think through what beliefs are reasonable, and we only want to keep track of strategies that are best responses to certain reasonable beliefs. So let's define this a bit more formally.

First of all, the key thing about this is these are strategies for player I. So here, as I said, I'm going to use some notation that maybe we haven't done before. I'm defining a set here.

And the way I'm building up the set using set builder notation is I'm going to say this set contains every strategy little si in the set big Si that satisfies some additional property. So this is kind of the universe of things I'm considering. And then I'm going to impose some constraints over here that narrow down which elements of this set I actually allow for.

So what does it mean for Si to be in here? Well, a strategy Si is in this set if Si is a best response to some belief, beta negative i. Now if I just stop there, we'd be back to what we did last week. We would just be thinking about being a best response to some belief.

But now I want to impose the restriction on this belief. I want to say it's a best response to a belief that only puts positive probability on strategies in this set, not any strategy. So on some belief beta negative i is satisfying.

Well, what's a way of saying this? Well, this belief assigns some probability to the opponent's strategies being in this set. And I'm going to require that that probability is 1. So beta negative i of S negative i prime equals 1.

So it's not enough for a strategy to be a best response to some belief. It has to be a best response to a particular kind of belief, a belief that maybe puts-- it could put probabilities on different strategies in this set, but it can't put any probability on strategies that are outside this set.

So the interpretation of this is if I know that my opponents are only going to play strategies in this set, what strategies for me are reasonable? And it's exactly the strategies that are best responses to a belief that doesn't put any probability outside that set, or equivalently, puts all the probability on the strategies in this set.

Now, we said from last week that there are two ways of thinking about this. We can think about what I should play and what I shouldn't play. So now we want to look at the other side of the coin, which is, what are the strategies that are strictly dominated by some mixed strategy? So let's go here. So what this stands for is strictly conditionally dominated.

So last week we talked about what it means for a strategy to be strictly dominated by another strategy. And it means some other strategy does better however my opponent's play.

But now we want to be a bit more permissive. Again, we're going to have the interpretation that I know that my opponent's strategies lie in here. So if I know that my opponent's strategies are in here, for dominance, I'm not going to require a strategy to do better however my opponent's play. It's going to be enough if that strategy does better however my opponent's play in this restricted set.

And that's what the conditional dominance means. It means conditional on my opponent's playing a strategy in this set. Conditional dominance means one strategy always does strictly better than the other strategy. So let me write it in words, and then we'll be a bit more formal. So again, this is a collection of strategies si and big Si, such that Si is strictly conditionally dominated.

Now, if I say conditionally dominated, strictly conditionally dominated, I have to say what it's conditional on. So strictly conditionally dominated I'll say given S negative i prime. So that's my condition. The condition is I know my opponents are going to play something in S negative i prime, and I'm interested in conditional dominance given that condition.

So it's strictly conditionally dominated given S negative i by what? By some mixed strategy. And I think it's often tricky to get-- it's easy to get confused here. Is it dominated by a mixed strategy of player i's or player i's opponents? What should come here? Yeah?

STUDENT: Player i.

IAN BALL: Player i, right? So it's always player i's strategies that dominate player I's other strategies. The beliefs are about the opponent's strategies, but the dominance relation is always between strategies by the same player, because the player is only comparing their own strategies. So by some mixed strategy sigma negative i.

So here's the definition, or here's the definition of this set. But I've of cheated a bit because I've used a new word that I haven't yet defined. So now I need to define that new word. What does it mean for Si to be strictly conditionally dominated given S negative i prime by some mixed strategy, sigma i?

Well, we already said it intuitively. It means that whatever strategy profile my opponents play in this set, sigma i is strictly better for me as player than strategy Si. And now we just need to translate this intuitive idea into math. So I'm comparing, as player i, what I get from sigma i to what I get from Si.

And again here, we're using our lowercase u's even though what I'm plugging in here is not a pure strategy but a mixed strategy. But remember what we said at the beginning, the interpretation is that this is just my expected utility over the lotteries induced by this mixed strategy.

So we're comparing sigma i to Si. We want Si to be dominated by sigma i. So we know the inequality has to go this way. And it's strict dominance, so it has to be a strict inequality.

But we're not going to require this for every strategy of my opponents. We're only going to require it for the strategies that satisfy the condition. So this is going to be S negative i S negative i, but for which S negative i for all S negative i that are elements of S negative i prime.

Now if there's two players this is pretty easy. S negative i is just a strategy of my opponent. But if there's more than two players, remember, say there's three players, then S negative i-- let's say i is 1. S negative i specifies what strategy player two chooses and what strategy player three chooses.

And in this case, S negative i prime is a collection of strategy pairs chosen by players two and player three. They're still choosing separately, but we're just putting them together mathematically into a vector.

Are there any questions on this? I think this is maybe the heaviest definition we've done. Any questions? OK, great. So now let's try to-- maybe I'll move to a new board. Let's try to recall what our theorem said last week in this new notation.

So often when we introduce new notation, it becomes a lot easier to express ideas that we've already expressed. So let's recall the theorem from last week. Well, what we said is that a strategy is either a best response or it's strictly dominated. It's not both.

But remember, last week we didn't worry about strict conditional dominance. We didn't worry about restricting beliefs. So it was a special case where S negative i prime was just S negative i. There wasn't any restriction on these things.

So our theorem from last week basically said two things. It said Si equals BRi of S negative i union SCD i of S negative i. Let's try to go through this in words. And first, let's think about what the union means.

This is a set of strategies. This is a set of strategies. When I take the union, I'm thinking of the larger set that puts together things that are in either of those two sets. So it's like I have two bags of strategies, and I pour them together into one big bag that contains all of the strategies.

So in words this says, every strategy in Si is either a best response to some belief. Notice, here I'm not-- there's no prime here. So this is a best response to an unrestricted belief. It would only have a restriction if I chose as S prime a smaller set of strategies.

So another way of saying this is if we go to this definition, well, if S negative i prime is just S negative i, then the fact that beta negative i of S negative i equals 1 is not a restriction at all. Because every belief assigns probability 1 to the set of all strategies.

So every strategy is either a best response to some belief, or it's strictly dominated. And again, this means strictly conditionally dominated. But strictly conditionally dominated conditional on any possible strategy just reduces to strict dominance. So strict dominance is a special case of strict conditional dominance.

But this is only one part of the theorem. What else did our theorem say? I mean, this says that every strategy is in one of these two sets. But it said one more-- there's one more thing we know about these sets. Yeah?

STUDENT: Each strategy is exactly one of the sets.

IAN BALL: Exactly one. So what's missing here? What haven't we-- what does this allow for that we want to rule out?

STUDENT: The last four do not have either function.

IAN BALL: But what we know is that that can't happen. So what we also know is that the intersection is empty. That is, no strategy is in both of these sets. So this is our "at least one" statement. And now we're going to-- just using the notation we said last week, this statement says that every strategy is in at least one of these two sets. But our second part of our statement was that every strategy was in at most one.

What does it mean to be in at most one? It means you're not in both. So what that means is BRi of S negative i intersection SCD of S negative i is empty. And the math notation for that, I'll just write this.

So again, this is a collection of strategies. This is a collection of strategies. I should have an i here. And all this notation says in math is that no strategy is in both of these sets. The intersection of these two sets is empty.

It turns out, though, that this holds more generally. So last week we had the special case with S negative i, where we didn't put any restrictions on these beliefs. It turns out this holds more generally if we put primes here. So this holds for any subset. So let's work through this.

We started with S negative i. Now we're using a smaller set of strategies. So let's try to understand what happens to these two sets. If S negative i prime gets smaller, so I'm considering fewer strategies that my opponents might play, what happens to these two sets? So this best response set-- does this smaller or bigger as S negative i prime gets smaller?

STUDENT: [INAUDIBLE].

IAN BALL: Say again. I think someone had it. Yeah, smaller, right? So this gets smaller. Why? Well, I'm not allowed as many beliefs right. This is the set of all strategies that are best response to certain beliefs.

If I shrink the set S negative i prime, it says I'm narrowing down the strategies that my opponents might play, which narrows down the beliefs I might hold. So how many strategies do I have that are best responses to this smaller, narrower set of beliefs? Fewer strategies. So this set is going to get smaller.

What about this set? As S negative i prime gets smaller, what's going to happen to this set? It's got to get bigger. Well, one way-- I'll answer in a second-- One way to see this is well, the union is staying the same size. So if one gets smaller, the other has to get bigger. But let's think through it mathematically.

A strategy is dominated by another strategy if that other strategy is just better, no matter what happens in here. It gets easier to be dominated if there are fewer strategies that my opponents could be playing. So it's easier to be conditionally dominated if the condition is stronger.

One way of thinking about it is just counting the inequalities. If S negative i prime is a really big set, we have a lot of inequalities here. These inequalities have to hold for every element of this big set. If Si prime becomes a smaller set, there are fewer things to check. It's easier to be dominated conditional on a smaller set of strategies. And therefore, the set of things that are dominated gets larger. That's what we see here.

And in the extreme case-- maybe one thing to think through. In the extreme case, what if this just has a single element? So special case. So what if S negative i prime contains a single strategy profile S negative i prime. So this is big S negative i prime. This is little s negative i prime. Then what happens to this set? What is this set going to be in words? Yeah?

STUDENT: Maybe [INAUDIBLE] be dominant.

IAN BALL: Not quite. So it's actually just going to be the strategies that are best responses to this strategy. So remember, before we got into all this fancy sets of strategies, we just defined what it is for one strategy to be a best response to another strategy. So if this set just has a single element in it, it just says this is the collection of strategies that are best responses to this particular strategy profile of my opponent.

But then this gets tricky. Then that means this is everything that's not a best response to this strategy. So how does that work out? Why is everything else strictly dominated then? That seems a bit tricky. That seems like we have a lot of strategies in there.

Remember, strict conditional dominance given one strategy is really simple. All it means to be in this set is that some other strategy does strictly better given this strategy of my opponents. But that's exactly the same as saying this strategy is not a best response.

If there's only one strategy of my opponent's, either a strategy is a best response to my opponent's strategy or it's not, meaning another strategy does strictly better given that strategy of my opponents. And that's exactly what strict conditional dominance means in this special case. That may have been a little quick. Sorry. Any questions here? Yes?

STUDENT: So is it always strict that every time we basically make Si-- when we shrink the S negative i prime, that the Ri S negative i prime is going to shrink? Or could it sometimes stay the same?

IAN BALL: It could sometimes stay the same. And, I mean, one way to see that is-- well, it's going to end eventually. So eventually, we're not going to be able to take things out. That's one way of seeing it. But you're right. We'll go through an example like this.

So it's not always the case that at each step things will shrink. But once nothing is thrown out, once everything stays the same for one round, then we're basically done with the algorithm. But we'll talk about that in a second. Yeah. OK. Any more questions about these relations here?

So now that we have our notation, we can formally define our algorithm. It took us a while. We started by motivating it. We gave these definitions. But with these definitions, we can write the algorithm in a really short clean way. So let's go over this.

So here's our algorithm. Maybe I'll write it kind in a more computer sciencey way. So what's our input to the algorithm? The input is just a strategic form game. And I'll be a little more precise-- let's say a finite strategic form game. And if I just write it like we did at the beginning, it's just a function u. And maybe I'll write S to Rn.

And in a sense, all of game theory is just studying functions like this. It's a function that says for every strategy profile the players choose, what are all the players payoffs? And this is a vector that tells me player one's payoff, player two's payoff, all the way up to player n's payoff.

And you can see this simple function gives me everything I need to know. Because the n up here tells me how many players there are. The S here tells me all the strategy profiles, and therefore, all the strategy sets for each of the n players. And then the function itself tells me all the payoffs. So that's really all I need to know.

Maybe I'll put a star by finite, and I'll say, what we often do throughout the course is when we talk about formal results, we assume things are finite to make sure nothing bad happens. The reason we do that is when you move away from the finite case, there are some weird cases that can go wrong. But then we'll discuss a lot of examples that aren't finite but where things are well behaved and nothing goes wrong.

So we will apply these ideas to the infinite case, but we just don't want to make any formal claims about the infinite case because you can come up with some really weird infinite discontinuous games where bad things happen.

But we will-- if on an exam I say, compute the rationalizable strategies in this game, you can't just say, oh, it's not a finite game. I'm not going to do it. You can't get out that easily. So we are going to apply these ideas to infinite games. But I'll formally define it here for the finite case.

So we're going to define things recursively. So let's start with our-- maybe I'll call it the base case. And the idea is we just want to go through the algorithm we did before. We want to start with all the strategies. And at each point, we want to throw out strategies that we don't think players will play. And we want to whittle down the set of strategies until we get to the good ones at the end.

So the base case-- we'll index things by k. So I is always going to index players. k will indicate the steps in this algorithm. So with k equals 0, well, we just say Si0 equals Si for all I equals 1 through n.

So we haven't done anything here. We've just said at step zero, what strategies are left? Well, the strategies that we started with. And that's true for each of the n players. So we've done nothing. So far, we've just formally gotten the algorithm going. We've initialized things.

And now let's look at the inductive step. And if you haven't taken an algorithms course or seen proofs by induction or seen recursive definitions, it just means after we start with k0, well, next we're going to go to step one. And then we're going to go to step two. And then we're going to go to step three. But it takes a long time to write out all those things.

So we're going to reason at an arbitrary step. We're going to say, let's say we've gone through steps up to k. How are we going to perform step k plus 1? And if you know how to go from k to k plus 1, well, we already know zero. So then we can go from 0 to 1, and then 1 to 2, and then 2 to 3, and so on.

So the inductive step-- well, what's our starting point? We're imagining that we've already defined these sets up to stage k. So what we have is we're given-- maybe I'll write it out-- S1k all the way up to Snk.

So these are the strategies for player 1 that have survived the rounds up to round k. I guess we're starting at 0, so it's really k plus 1 and them, but that rounds up to k. And Snk is the set of strategies for player n that have survived so far up to round k. And what we want to do is define which strategies are going to be left at round k plus 1.

So what do we need to do? We're going to define maybe here, say, for each i, Sik plus 1. So before I write down the formula, let's just understand what we have. We have S1k through Snk. We want to define Sik plus 1 for every player I. So Sik plus 1 is the set of strategies for player i that survive up to round k plus 1.

Any ideas based on the notation we've done given so far what could be here? I mean, intuitively, we want to keep in all the strategies that are best responses to some belief about our opponents that's consistent with what we've calculated so far.

So what we know at stage k is that it's only reasonable for player one to play something in here, all the way up for player n to play something in here. So given that knowledge that player i has, what is reasonable for player i to play at this round of the algorithm? Any thoughts? Yeah?

STUDENT: [INAUDIBLE] i of S minus the other--

IAN BALL: Minus i, yeah.

STUDENT: To the k.

IAN BALL: Exactly right. So this is going to be BRi of S negative i to the k. Now let me first define what this is. So I think we know this notation. But let's be clear. This is something like S1k all the way up to Si minus 1k times Si plus 1k Snk.

I think that makes it seem fancier than it is. All it means is I know the strategy set at the k-th round for each of my opponents, and I just want to put all those together, making sure to skip myself. Because I'm player i, I don't need to form beliefs about my own strategies.

So this is-- you can think of it as-- or it exactly is the set of strategy profiles of my opponent, such that each player's strategy in that profile lies in the corresponding strategy set that we got at the end of round k. And more mathematically, from our definition here, these are the set of strategies for player i that are best responses to some belief player i could hold that only puts positive probability on strategies that are in this special restricted set.

Now, there's another definition we could use using our theorem over here. How else could we define this step that would give us the same answer? Yeah?

STUDENT: We could define the set as basically Si, and you take out the [INAUDIBLE] SCD i S negative i to the k.

IAN BALL: Exactly. So just like we always said, there's two ways we can do it. We can keep in what's good, or we can throw out what's bad. So we can also define it as everything in Si except-- so when I use this set minus thing, I want to say the set of things-- so everything that's in here, but not in here. And that's SCD i of S negative ik. And our theorem over here exactly tells us that those two are the same.

So you could always use this definition. You could always use this definition. It's not going to make a difference. This definition is maybe more consistent with the name "iterated elimination of strictly dominated strategies." That's literally what we're doing here.

But we could also say iterated retention of best response strategies. And that's basically what we're doing here. Though, I just made that up. That's not a name that's used. But that's one thing we could do. Great. So we're not done yet because all I've told you is how to go through each step. Step-- oh, Yes

STUDENT: Why isn't the S sub i S sub i to the k?

IAN BALL: Here? That's a good point. It turns out you could do that, and you get the same answer. So in the infinite case, you have to be a little careful. But in the finite case, you're right. If you put Si to the k here, you would get the same answer.

I think conceptually, I like to think-- I think this is conceptually more-- better captures it, that here I am. I'm player i. All I know at this stage is that this is what the players could do, and I want to throw out things that are strictly dominated. So it doesn't-- there's no reason for me to really think back to what I did before. So I think conceptually maybe this is more in the spirit of the algorithm.

But you're exactly right, that mathematically this is going to be a subset of Sik. And if I put Sik here, and often when we do the algorithm, that's kind of what we implicitly do in our heads. So you could put the k here, and it wouldn't-- in the finite case, it wouldn't change things. You have to be a bit careful in the infinite case. Yeah. Great.

So we've defined these steps. But in general, this could go on forever. So if we wanted to be really formal, then what is the output of the algorithm? Well the output of the algorithm is we think of which strategies remain after we do this infinitely many times. It's a little weird to think about. But the output technically we might call Si infinity. And this is the set of strategies that survive every single round.

So if we want to be really mathy about it, we would say this is Si0, intersect Si1, intersect Si2, intersect, and dot, dot, dot. And we could write that as the intersection from k equals 0 to infinity of Sik.

So formally what does the set mean? It means I only include a strategy in this set if it survives round zero, it survives round one, it survives round two, it survives round three, and on forever. And we call this the set of rationalizable strategies for player i.

So if you wanted to be really precise, we might say that iterated elimination of strictly dominated strategies is an algorithm or a procedure, and that the output of that procedure is the set of rationalizable strategies. You might say rationalizabilty is a property of a strategy. And IESDS is a procedure that we go through in order to identify the strategies that have that property. But on exams, we're not going to be too picky about exactly this distinction between the algorithm and the output of the algorithm.

But you might say, well, wait. We have no hope of doing this. We have to go through infinitely many steps. That's a mess. So fortunately, when we're in a finite game, the algorithm is always going to end after finitely many steps. And what do I mean by that? So I'll say fortunately, this ends after finitely many steps.

Now, what does ending mean? It doesn't mean I only have one strategy left for each player. That happened in our example, but that doesn't always happen. But what ending means, kind of going back to the question earlier, is I get to a point where I can't throw out any more strategies.

And if I don't throw any more out at one round, then I'm not going to be able to throw any more out at the next round. And basically, the algorithm is just going to become static. There's going to be nothing else that happens.

So formally, what it means is for some big K-- so after big K steps, subsequent steps don't change anything. So I'll say for some big K, SiK-- that's a big K-- equals SiK plus 1 equals SiK plus 2, and so on for all i.

So what that says is once I get to stage K, if I were to apply another step of the algorithm, nothing would change. And if I were to apply one more step, nothing would change.

How do you know you've gotten to that point? Well, as long as you get to a step where nothing is thrown out, you're done. But let's be really careful. It's not enough for there to be one player where you throw nothing out. If you throw strategies out for any player, you might have to keep going. So the algorithm is only done when every single player's strategy set stays the same from one round to the next.

You can even see in our example here, when we went from step one to step two, player two's strategy set-- sorry-- player one's strategy set did not change. But the algorithm wasn't done because player two's strategy set did change. And for that reason, we had to do one more step. So it's only when everyone's strategy set is fixed that you're done. So keep going until nothing changes, and then you're done. And what's left are the rationalizable strategies.

Now, let me give a tip, I think, for exams, for problem sets when you're actually doing this on your own. This algorithm has a-- or before I do the tip, let me say, how do we know that it has to end if we want to be a little mathy and think, wait. How can we be sure it ends? Why can't it go on forever? Anyone know?

Well, there's only finitely many strategies to begin with. And at each round, we just throw strategies out. We never get bigger. So if we have finitely many things, and we keep throwing something out, we can't throw things out forever. Eventually, we can't get any smaller. Yes?

STUDENT: Does that mean you can never swap strategies? Like, [INAUDIBLE] like, throw one out [INAUDIBLE].

IAN BALL: Exactly. That will never happen. Yeah. So at each step, these sets are getting smaller and smaller and smaller. And this exactly goes to the property we were describing about how the set of best responses gets smaller as the set of opponent strategies gets smaller.

So at each round, everyone else's strategy set gets smaller. But that means the set of best responses I have gets smaller, which means the set of best responses they have gets smaller. And at each stage, we're just throwing things out.

Another way to see-- well, yeah. I guess it's not immediately obvious from this. But again, this set gets smaller, and this set gets bigger. We throw out more things, and we keep fewer things.

If we didn't have that property, then it could go forever. We'd worry that we might get in cycles where we take a strategy out, put it back in, take it out, put it back in. That can't happen. There's no cycles. We just get smaller and smaller until we stop. Yes?

STUDENT: So when you say it's a finite strategic form game, does that mean the number of players is finite, as is the number of strategies for a player?

IAN BALL: Yes. So n is finite. I mean, we're never really going to think of n infinite. But technically, it means n is finite. But the real substantive content of it is that the Si's are also finite. Yeah.

So here's my tip for doing this. Because of this property that things get smaller at every stage, if you throw something out that you're not supposed to, you're done. You're really in trouble. You've messed up. There's no coming back from it. But if you don't throw something out that you're supposed to throw out, you can always throw it out next time.

So it's always safer to keep things in because once you throw it out, it can't come back. But if you keep it in, well, maybe it just stays in a little too long, and you'll throw it out next round. So it's better to err on the side of inclusion.

And for that reason, my suggestion would be when you're going through it, think about dominance. So I would say use SCD during the algorithm. Why?

Well, if you're looking for strategies that are strictly dominated and you miss one, that means you forget to throw something out, and you can throw it out next time. If you look for strategies that are best responses and you miss one, then you're going to throw that strategy out, and you can never recover from that.

Now, the only problem, though, is at the end, if I'm using this definition, there's a risk that I might have kept in too many strategies at the end. So my advice would be at the last stage, once you think you're done, then use this definition, and check that everything that's left is a best response to some belief about your opponents.

So use SCD during the algorithm, and then use BR at the end. And this is a safe way to make sure you never throw anything out you shouldn't throughout the algorithm. But then at the very end, check that there's nothing left that you should have thrown out.

So that would be my advice for how you go about it if you want to be safe. Maybe it's a bit slower, but I think that's a safe way to do it. Any questions on that? Great.

So let's now do a few examples. So another simple example would be the prisoner's dilemma-- so cooperate defect, cooperate defect. I think last time we may have used the phrases m and f, mum and fink. But in the prisoner's dilemma it's also common to use cooperate, defect, cooperate, defect.

Remember, cooperate means cooperate with the other prisoner, not with the authorities. That can be a little confusing. So cooperation is a property of what the players in the game are doing. So it doesn't mean testify against the other prisoner. It means cooperate with the other prisoner by not testifying.

So here in the prisoner's dilemma, we have-- let's go through the algorithm. Well, first, what is S10? Well, this is just the set of strategies for player one. We already see what it is. It's C, D. So that's easy. And then S20 also equals C, D.

So now let's go-- so this is step k equals 0. Here's k here. Now let's go to step one. So in step one, what happens? Let's follow our tip, and let's only throw out strictly dominated strategies. So what do we throw out here?

Well, cooperate is strictly dominated by defect. If my opponent cooperates, I do strictly better by defecting. And if my opponent defects, I also do strictly better by defecting. So D strictly dominates C for each player. That means I throw out D.

And what do I keep in? Well, C. Sorry. I throw out C, and I keep in D. Let me say that again. I may have jumbled my words. C is strictly dominated by D, so I throw out C and keep D.

And now I could go one more stage. But actually, here one way you always know you're done is if each player only has one strategy left. But let's make sure we didn't make a mistake. Let's do our best response thing at check at the end. Is D a best response to D? Well, we can actually go through our procedure here.

So I'm just doing the underlining that we did before. And sure enough D is a best response to D for both players. And we know-- well, we don't know we've done it right but we at least know we didn't obviously do it wrong. So we're done there.

Let's do one more example. Maybe I'll call top, bottom, left to right. So again, two players, two strategies each. Let's do it like this, and let's go through our procedure.

So again, step zero is easy. We're just writing down the strategies the players have. Player one strategies are T and B. Player two strategies are L and R. So we have S10 equals TB, and S20 equals L.

And I warn you when you look at this, it gets a little confusing when you see things like S11. But I'm always going to use the convention that the subscript is the player, the thing at the bottom, and the superscript is the round number, or the stage of elimination.

So I'll call these things k at the top and i at the bottom. But if they take actual numbers, just remember the bottom is the player, and the top is the stage that we're at. So that's step one-- step zero.

Now, let's go to step one. What do we throw out here? Yeah?

STUDENT: Can a strategy [INAUDIBLE].

IAN BALL: Right. So we don't throw anything out. And how can we see that? Well, let's go through our best response thing. If I'm player one, if my opponent plays L, I'd rather play T.

If my opponent plays R-- we haven't gotten to a case like this. But here, I have multiple best responses. And it's really crucial when you have multiple best responses to underline both of them because they're both reasonable ways you could respond.

So given my opponent plays R, I get 0 from both T and B So I'm going to underline both of them. And then we'll do the same thing the other way. If player one plays T, player two wants to play L. If player one plays B, player two is indifferent between L and R, so we underline both of these.

And if we look at our underlining, we see that T is indeed a best response for player one. It's a best response to L, and actually also to R. But L is enough.

What about B? Is B a best response to anything? Yes. B is a best response to R. It's underlined here. So both T and B are best responses to some belief. And similarly, both L and R are best responses to some belief.

L is a best response to both T and B because we can see L underlined twice. And R is a best response to B, but not to T. But that's OK. It's a best response to something. So we get S11 equals TB, and S21 equals LR. So do we have to keep going, or are we done? So now we're done.

But the reason we're done is that player one's strategy set didn't change, and player two's strategy set didn't change. But remember, we're only done because every single player's strategy set stayed the same. It's not enough for just one to stay the same. So we're done here. We're done here.

And these two examples illustrate that, in some cases, this rationalizability procedure gives us a very strong prediction about what we'll play in the game. But in other cases, it doesn't really give us any sharp prediction at all. We're not able to rule out any strategies using this concept of rationalizability alone. Let me check the time. Great.

So let me do one maybe more substantive example to finish that actually has continuous strategies. So let's go back to the beauty contest game that we played on the first day of class. And I think some people were starting to go through this kind of reasoning. Here, we can formally solve this game.

So let's go to the beauty contest. But this time, it's a bit easier to say that you get to choose a number, any number, any real number between 0 and 100. So you're not restricted to integers here. So each player gets to choose.

So we have players i equals 1 through n. And each player i gets to choose a number between 0 and 100. So I'm following the notes here and using x to be the number that you choose instead of s.

Let's write down our payoffs. So let's say I'm player i, and the numbers chosen are x1, x2, all the way to xn. So we're going to consider here the alpha beauty contest. So we're going to have some number alpha that's between 0 and 1. And your goal is for your guess to match alpha times the average of everyone's guess.

So on the first day of class, we played. And the second day, we played the beauty contest with alpha equals 2/3, I think. And to make things a bit easier, instead of just saying you either win or you don't, we're going to say that your payoff depends on how close you are to alpha times the average.

You get the highest possible payoff if you're exactly alpha times the average. And the farther your guess is away from alpha times the average, the lower your payoff is. And we're going to specifically do that with what's called a quadratic loss function. So you lose the square of your distance from your ideal guess.

So what is that going to be? It's going to be minus something squared. So the thing squared. This is your loss, and you're subtracting this loss. And what are we going to square? We want to square the distance or the difference between what you actually guess and the ideal guess, which is alpha times the average of the class.

So this is player i's payoff. So I have xi here. That's what I actually guess. What's the best possible guess? It's alpha times the average of everyone else's guess-- or average of everyone's guess, I should say. So that's minus alpha times x1 to xn over n. And this is kind of a subtle aspect of this game, that my own guess contributes to the class average.

So here, this is the class average. There's n people in the class. Each of them guesses a number. So I divide by n to get the average guess in the class. Then I multiply the average guess, which remember, I is player i contribute to. I multiply that average guess by alpha. That's the ideal guess. That would be the best possible guess I could make.

If I guess exactly that, then my payoff is zero. That's great. But the farther my guess is away from my bliss point, or the ideal guess, the greater my losses, or the more negative my payoff is.

It turns out because-- it's a little tricky to analyze this directly because when I think about my guess, I'm also changing the class average at the same time. So it's better to just factor out the xi. So let's write this a bit differently.

So we have xi here. But then, there's also an xi here, and it's multiplied by alpha over n. So what I get is 1 minus alpha over n times xi. The 1 times xi comes from here, and the minus alpha over n times xi comes from over here.

And then what I'm left with is alpha times the sum of everyone else's guess divided by n. But that's kind of awkward to deal with because there's only n minus 1 people, and I'm dividing by n So I'd rather divide by n minus 1. So let's write this as alpha times n minus 1 over n times x negative i bar, where this is the sum over everyone but me divided by n minus 1.

So all I've done is multiplied and divided by n minus 1. It's a little algebra here. But I divided by n minus 1 here. I wasn't really supposed to do that. So I undid it by multiplying by n minus 1 over here. And I still have the n that I'm supposed to have on the bottom.

And this is just nicer notation because this tells me, what is the average of everyone else's guess? And that's really the way you think about this problem. You think about your guess, and you think about the average of everyone else's guess.

It's a bit weird to think of the class average, which partly reflects your guess and partly reflects everyone else's guess. So this is just a slightly nicer way of writing it, though it makes things a little messier. Let's go down here.

So from this formula, it's also a bit easier to calculate my best response. So let's imagine, hypothetically, that I knew what the average of everyone else's guess was. Then I would want to choose my xi so that this equals this. And this should be squared.

I want my loss as small as possible. I want this expression to be 0. So I want this to equal that. And now I'm going to solve for xi. So I want xi to be this divided by this. Just a little bit of algebra, which you can work through at home if you'd like.

And what you'll find is that the kind of best response, if I solve for this, is I get xi equals-- well, I get the right side divided by 1 minus alpha over n. And the n's cancel nicely. So I get alpha times n minus 1 over n minus alpha times x negative i bar.

So what this says is if I knew that this is what everyone else was going to guess, or more specifically, this is the average of everyone else's guess, then this is my best response. This is my best guess.

Now, at first you might have thought you would want to guess alpha times what everyone else guesses, but that's not what we get. Why do we get this kind of messy term here?

Well, first, let's look at this number. Is this number bigger than 1 or less than 1? Let's assume n is greater than 1. Otherwise, this is not a very interesting game. This is less than 1 because I'm subtracting 1 on the top, and I'm subtracting something smaller on the bottom. So that means the bottom is bigger than the top. So this is less than 1.

So that means it's best for me to guess not alpha times the average of everyone else, but strictly less than that. So this is less than alpha x minus i bar. Does anyone have any intuition? It's a bit strange. If you thought about this game, you might think you just want to guess exactly alpha times everyone else's guess. Yeah?

STUDENT: Because your guess is also going to be part of the average, and you're bringing it down.

IAN BALL: Exactly right. So if I said OK, everyone else, the average is x negative i bar. Let's say I guessed alpha times x negative i bar. Well, that sounds good. I'm guessing alpha times the average. But now the average is strictly smaller because my guess of alpha times the average is smaller than the average.

So exactly as you said, I've pulled down the class average, and now I want to guess a little bit lower. And that's exactly what this is reflecting, the fact that my guess directly affects the class average. And when I'm trying to guess a fraction of the class average, my guess is always going to pull it down.

But not so important. Let's just give this a name. We'll call this alpha n. So technically it depends on this parameter, alpha and the number of players we have. And it's the number of players grows large, it's going to get really close to alpha because your effect on the average is going to become very small.

But now let's try to go through our algorithm in this game. OK, running out of time, so let's just do it quickly. Si0 equals 0 to 100. Si1-- well, the highest my opponents could guess would be they all guessed 100. And in that case, I'd want to guess alpha n times 100. The lowest I'd want to guess is 0, if everyone else guesses 0.

So we're going to get 0 to 100 alpha n. And you can see the pattern. Si to the k is going to be 0 times 100 alpha n to the k. Because each time, if I know that my opponents are only guessing up to 100 times alpha n, I want to take that and multiply it by alpha n again.

So in the limit, what we find is that Si infinity equals 0 for all i. So this game is dominant solvable. And the unique rationalizable strategy for every player is to play exactly 0. That was a bit quick, but a variant of this is on the homework, so I didn't want to make it too easy for the homework. So you'll get to work on the problem set.

OK, I'll see everyone next week.