[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: So today the topic is zero-sum games. And what does this mean? A zero-sum game is a game that features exact conflict of interest. So maybe I'll call it a complete conflict of interest.

Another way of saying this is that one player's gain is the other player's loss. And whenever I'm talking about losses and gains, remember, this is about utilities, not about money. So it could be about money if we're risk neutral, but more specifically, one player's utility gain is the other player's utility loss.

So what are some examples where this is a realistic assumption? A lot of these examples are actually more parlor games, games like chess and checkers. So if we think about, say, chess or checkers, what happens?

Well, assuming that people only care about the winner of the game-- you could imagine someone who just loves playing chess. They want the game to last as long as possible, or they want a beautiful game of chess. But we're thinking about people who just care about winning and losing.

Well, what are the possible outcomes? Well, either player 1 wins, and if player 1 wins, we'll think of payoffs as being 1 and negative 1, 1 for the winner and negative 1 for the loser, player 2.

The other option is player 2 wins. And then, of course, the payoffs are the other way. Negative 1 for player 1 who lost and 1 for player 2 who won, and then the last possible outcome is a draw, a tie. And in that case, we'll say the payoffs are 0, 0.

So here, you can see that the player's interests are exactly opposed. In order for one player's payoff to go up, the other player's payoff has to go down. And in this case, the sum of the player's payoffs is always 0. So we see that payoffs sum to 0.

Another common example of this would be-- I should start in the outer-- games with monetary trade or monetary exchange, so maybe money exchange. But there, we have to be careful if we're playing-- if we're bargaining about how much money I give to you or how much money you give to me, this can still be a zero-sum game, but we need two assumptions. So I'll say under two conditions.

Well, the first condition is that no money enters or leaves the game. If we're negotiating and then someone's parents come along and give us both money, that's not a zero-sum game anymore. So there's no money in or out.

And then what's the second assumption we need in order for this to actually be a zero-sum game? Thoughts? Well, it's important that we're risk neutral.

Because if we're risk averse, for instance, and the terms of our trade were that we're going to flip a coin, and someone's going to win a lot of money, and someone's going to lose a lot of money, well, if we're both risk averse, that bet might be bad for both of us in expected utility because we both have a chance of losing a lot of money.

So the fact that no money comes in or out doesn't necessarily mean that no utility comes in or out unless we make the further assumption that players are risk neutral. But under these two assumptions, money exchange is also going to be an example. Now, today we're going backwards historically.

So zero-sum games were some of the first games that were studied in game theory. They were studied by mathematicians who were motivated by situations like chess and checkers. Subsequently, economists started getting interested in game theory.

And in a lot of economic situations outside of games, outside of parlor games, the situation is not zero-sum. And you often hear people in politics talking about this, saying, this is not a 0 sum game. There's some arrangement we can make that makes us both better off.

So I'll say, maybe, in the real world, things are often not zero-sum. Maybe I'll say rarely zero-sum. And that's why game theory has moved on a bit beyond zero-sum games and looked at things like Nash equilibrium that can capture non-zero-sum games.

But mathematically, zero-sum games have a very nice, elegant theory. So we're going to cover it today. And it does apply to a lot of games like chess and checkers and poker that are relevant to some people. Yes, question.

AUDIENCE: [INAUDIBLE] risk neutral, that just kind of means the money is taken at face value?

IAN BALL: Yes, so formally, risk neutral means that your von Neumann-Morgenstern utility function, as a function of money, is a linear function. So that you are indifferent between, say, getting $10 for sure and a lottery that pays you 0 with probability 1/2 and 20 with probability 1/2. Good question.

Any other questions? So I'll say over money maybe to be more clear. Question or no? No, OK.

All right. So in the real world, rarely zero-sum, but maybe I'll say something positive. But the theory is elegant. And if you play poker and you play these games, it can often be a valuable benchmark for some situations in the real world that are close to strictly competitive, close to complete conflicts of interest.

All right, so let's now formally define what we mean by a zero-sum game. So throughout today, when we talk about zero-sum games, we're always going to be focusing on finite two-player strategic-form games.

Finiteness is kind of a technical assumption. A lot of the results extend beyond finiteness. Two-player is really essential. Games that only have two players with opposed interests behave quite differently than games that have more than two players. So this is really a central assumption.

And then the fact that we focus on strategic forms is just a modeling choice. A lot of games like chess and checkers take place over time, and we would naturally model them in the extensive form. But remember, starting with the extensive form, we can always represent that game in the strategic form. And we're going to assume that we've already done that transformation.

So how do we represent these games where we have two payoff functions, player 1's payoff and player 2's payoff. So let's just remember, player 1's payoff depends on the strategy that player 1 chooses and the strategy that player 2 chooses. And similarly, player 2's payoff depends on both player's strategies. So this is just a two-player game.

What makes it zero-sum? Well, it's what you would think. The sum of the payoffs is 0. So this is zero-sum, what does this mean? It means that player 1's payoff from some strategy profile s plus player 2's payoff from the same strategy profile s is 0 for every strategy profile s.

And remember, this is a strategy profile. So this is something like S1, S2. So when I write it down it looks like just one equation, but really, we have many equations. We have a separate equation that must hold for any pair of strategies that the players play.

And now we're going to make a simplification. If this is the case, then we can just do a little algebra. And we see that that means that u2 of s is equal to negative u1 of s for every strategy profile s.

The payoff or utility that player 2 gets is just negative of the utility that player 1 gets. And what that means is that if we're in a zero-sum context, we don't really have two utility functions. We really only have one function, and that determines what the other person's utility is because, once we know one player's utility, we know that the other player's utility must always be the negative of that.

So we're going to adopt a convention. So here's our convention for zero-sum games. We're basically only going to work with player 1's payoff. So let's let u equals u1.

That's going to be player 1's payoff. And that means we have in mind that u2 is equal to negative u. But we're only going to keep track of this single-function u. So we're going to focus on player 1.

And this means that u is going to represent player 1's gain, but it will also represent player 2's loss because a high value of u means a large negative value for player 2, which is a loss, so this represents player 2's gain. So instead of thinking of each player-- loss, thank you.

So instead of thinking of each player as having their own utility function that they're trying to maximize, we now just have one function u, but now the players are doing different things with it. So player 1 is trying to maximize u. But player 2 is trying to minimize u.

And we're going to maintain this throughout the course. Now, this creates this asymmetry between the players. Player 1 is the maximizer. Player 2 is the minimizer. I find that zero-sum games, it's really easy to get yourself confused, so hopefully, I won't get confused today.

I don't think what we're doing today is the hardest material we're going to cover in the course. But I think it probably is the most confusing. It's the easiest lecture to get confused on. So if you have a question, please do ask, and hopefully, I won't confuse myself with all these mins and maxes. OK.

So that's going to be our convention. And now let's just go over one really simple example of a zero-sum game to try to get our bearings here. And the classic zero-sum game would be Rock paper scissors. So let's look at this.

Now, we're going to represent this in strategic form. So we're going to have a matrix here. Player 1 can choose rock, paper, or scissors. And player two can choose rock, paper, or scissors.

Now, in the past, when we represented a strategic-form game, every box in this matrix had two numbers. It had the payoff for player 1 and the payoff for player 2. But now we're going to follow our convention, and we're only going to put one number in every box. But that one number is going to represent the gain for player 1 and the loss for player 2.

So let's just fill this in cleverly. Well, if we go down the diagonal, in this case, everyone-- the two players are making the same move, and that means the game ends in a tie. So if the game ends in a tie, we'll say the payoff is 0 to both players. So along the diagonal, we have 0, 0, and 0.

Now we want to put a 1-- remember, this is player 1's payoff whenever player 1 wins the game. So hopefully, people remember the rules for Rock paper scissors. Let's see.

So scissors beat paper for player 1, so we have a 1 here. Paper beats rock, so we have a 1 here. And rock beats scissors, so we have a 1 here.

And in the other three spots, the winner is player 2, which means the payoff to player 1 is negative 1 because player 1 is the loser. So if player 1 plays scissors and player 2 plays rock, player 1 loses, similarly, player 1 loses. Player 1 loses.

So whenever we're talking about zero-sum games, you're going to see something like this. You're only going to have a single number in every box rather than two numbers. But we're still completely specifying the game because of our assumption that that is a zero-sum game.

So now the question is, how should someone play in this, maybe how to play? I think people probably know what's a good strategy in Rock, paper, scissors, but let's think it through. One way to think it through is let's take the perspective of player 1.

And one thing player 1 could think about is they could say, if I choose a particular move or a particular strategy, what's the worst that can happen? I'm going to say, what's the worst that can happen?

This is a very pessimistic or conservative way of playing, but it's one thing we can think about. So let's say I play rock and I'm player 1, what's the worst thing that can happen? Well, player 2 could play paper, and they would beat me. So the worst thing, I'll put an arrow here. The worst thing is player 2 plays paper.

And my payoff, therefore, is negative 1. Negative 1 is my payoff. Similarly, if I play paper, again, the worst thing is pretty bad. If I play paper, then my opponent could play scissors.

And then, again, my payoff is negative 1. And symmetrically, if I play scissors, my opponent could play rock, and I could get a payoff of negative 1. So if I'm worried about the worst thing that can happen and I only look at these pure strategies, it looks pretty dire. Whatever move I make, the worst thing that can happen is the other player plays the one move that beats me, I lose, and therefore I get a payoff of negative 1.

But what if I instead play a mixed strategy? Is there any mixed strategy that I could play that would ensure that, even in the worst case, I do pretty well? Any thoughts? Yeah.

AUDIENCE: [INAUDIBLE] [? randomize ?] shows that you play [INAUDIBLE].

IAN BALL: Great, so let's look at the mixed strategy. I'll write it down here, and I'll call this strategy 1/3, 1/3, 1/3. And the meaning of this is I play rock with probability 1/3, paper with probability 1/3, and scissors with probability 1/3.

And it's not too hard to see that if I play this strategy, then whatever my opponent does, I'm going to win a third of the time, I'm going to lose a third of the time, and I'm going to tie a third of the time. And therefore, my expected payoff is going to be 0. So if I fill in with this row what my payoff is when I, as player 1, use this mixed strategy. And my opponent uses each of these strategies, we're going to get 0, 0, 0.

And the mathematical way to do this is we're literally taking a third of this row plus a third of this row plus a third of this row, and when we add those rows together, we're just going to get 0 at the bottom. So here, what's the worst case, the worst thing that can happen? Well, actually, in this case, it doesn't really matter what my opponent does. Whatever my opponent does, I get a payoff of 0.

So this suggests, and if we carefully went through it, we could see that no other strategy is going to do better in the worst case. There's going to be no strategy in Rock paper scissors that ensures that I always get a positive payout. So this is actually going to be a very good strategy if I'm concerned about the worst thing that can happen.

And I don't think it's really obvious why I should be worried about the worst thing that can happen. I mean, it's reasonable, but there are other reasonable things I can do. But I promise you that it's going to turn out that this turns out to be a very useful and powerful way to play.

So we're going to introduce some formal terminology about this kind of reasoning, with the promise that later on it will be very fruitful. So let's think about this worst case or reasoning. So let's formally think about what is-- maybe I'll call this worst-case reasoning.

Another way of thinking about this is maybe safe or conservative or secure play. What I want is I want to find a way to play in the game so that, whatever my opponent does, I'm guaranteed to do pretty well.

So I'm interested in what payoff I can guarantee, even In the worst case, no matter what my opponent does. So let's introduce this idea formally in a more general case beyond Rock paper scissors. So we'll think about player 1, so for player 1, they're looking at gains. So let me put player 1 here, and we have player 2 here.

So when player 1 thinks about the worst case that can happen, they're saying, what is the worst gain I can get as a function of my strategy sigma 1. Now, mathematically, this is just saying, well, the worst thing that can happen is that my opponent plays a strategy that makes my gain as small as possible. So the worst case, I'm going to take a minimum over sigma 2. Maybe I'll put-- remember our notation here, that capital sigma is the set of mixed strategies for player 2. And I'll write u of sigma 1, sigma 2.

So let's think through this really carefully. We're saying I'm player 1. Suppose I use the mixed strategy sigma 1. What is the worst thing that can happen?

Well, the payoff I get is going to depend on the strategy sigma 2 that my opponent plays. But I'm going to evaluate a minimum to think of the worst thing that could possibly happen, and I'm going to call that my Worst Gain, WG, worst gain from sigma 1.

And gain reflects the fact that, for me, u, the payoff, is a gain for me. It's a good thing for me. One observation here that we often make is that I don't actually have to minimize over all mixed strategies. If I just minimized over pure strategies, I'd get the same answer. So we could also write this as the min over s2 and s2 of u of sigma 1, s2.

And the reason is that my payoff from a mix, if my opponent plays a mixed strategy, is just the average, an average over the payoffs I get if my opponent plays pure strategies. So the worst thing that can happen will always actually be a pure strategy. Or there will be a pure strategy that's worst.

And now let's do things symmetrically for player 2. We want to do-- think about this secure worst-case reasoning for player 2. But now player 2 is focused on their loss. So I'm going to call this WL of sigma 2 because, remember, these numbers in the matrix for player 2 correspond to losses to player 2. So I want to say if I'm player 2 and I choose a mixed-strategy sigma 2, what is the worst possible thing that can happen?

The worst thing that can happen is if my loss is as large as possible. So I'm now going to write here a maximum over sigma 1 and sigma 1 of u, sigma 1, sigma 2. This says how much I lose depends on what player 1 does. And the worst possible loss is if player 1 chooses the strategy that makes my loss as large as possible. So we have a maximum here and a minimum here.

And again, this is the same as max over s1 and s1 sigma 2. And we might call this my worst loss. So just to check our understanding, if we went back to this example, if I'm player 1, if I played sigma 1 equals R, what would be my worst gain from R In this game over here?

It would be negative 1 because if I play R, the worst thing that can happen is my opponent plays paper, and then I lose, and I get a payoff of negative 1. On the other hand, the worst gain from the mixed strategy a third, a third, a third is going to be 0. That's what we showed down here.

And similarly, for player 2, the worst loss-- and maybe since we haven't worked with player two much, let's be careful. The worst loss for player 2, if they play R in this game would be what? Let's say I'm player 2 and I play rock in Rock paper scissors, what is my worst loss?

So we have to be careful here. It's actually 1. You might be tempted to say negative 1, but remember, this is how much I'm losing. So the worst case for me, if I'm player 2 and I play R, well, I'm trying to minimize how much I lose.

So the worst loss is if my opponent plays paper and I lose utility 1, my loss is 1. So just keep this in mind, and this is why things can be a little bit confusing here. Yes.

AUDIENCE: I understand structurally why a mixed strategy is the same as pure strategy for those formulations that you have there. I want to confirm, does this just hold for this scheme? Or is it a general result for zero-sum games?

IAN BALL: So it's a very general result that, once I fix one player's strategy in a two-player game, and now I want to maximize the payoff of the other player or minimize the payoff of the other player, that maximum or minimum will be achieved by a pure strategy. More generally, in an n-player game, if I fix the strategies of all but one of the players, n minus 1 of the players, and now I contemplate a choice of strategy for that last player, that payoff will be maximized or minimized by a pure strategy, maybe also by a mixed strategy, but in particular by a pure strategy.

But we have to be really careful here. On the other hand, it's not true that if I take a step back and look at the worst case gain, that my best worst-case gain can be achieved by a pure strategy. That's not true. And that's exactly what we saw over here.

Because if I restricted myself to a pure strategy, then my worst-case gain was always negative 1. And it was very important that we allowed myself to use mixed strategies. So mathematically, what's going on is that once I fix one player's strategy, the utility function as a function of the other player's strategy is a linear function. And a linear function over probabilities achieves its maximum at a corner solution at the endpoints.

But the worst-case gain function, WG, is no longer a linear function because it's a minimum over linear functions. It's actually going to be a concave function, and that's mathematically what's going on. Good question. Any other questions on this?

So we've defined, we've thought through worst-case reasoning. And now we want to go a step further and talk about something called security strategies, which just formalizes what we talked about in this game before.

So what is a security strategy? Well, when I'm choosing what strategy to play as player 1, I evaluate each strategy according to its worst-case payoff. And I choose the strategy that maximizes that worst-case payoff. So for player 1, we say a strategy sigma 1 star is a security strategy.

If the worst-case gain from sigma 1 star is better than the worst-case gain I would get from any alternative strategy.

So I'm comparing if I use strategy sigma 1 star, what is the worst-case gain I get? If I use a different strategy, sigma 1, what is the worst-case gain that I get? And I would like the worst-case gain to be as large as possible. And that's what it means to be a security strategy.

And then for player 2, things are just symmetric. For player 2, we say sigma 2 star is a Security Strategy. I'll just say SS.

Let's be careful here. Now we're worried about the worst-case loss of sigma 2 star. And for player 2, they're minimizing. They care about losses, so sigma 2 star is best for them if the loss is as small as possible.

So they want the worst-case loss from sigma 2 star to be less than or equal to the worst-case loss from sigma 2 for all sigma 2 in sigma 2. And here, going back to the question over here, it is really important that we contemplate mixed strategies here. We can't just look at pure strategies. We have this definition.

So it's nice to have some notation here. If we think about player 1-- come down to this.

So now that we have this notion of security strategies, we can start thinking about payoffs. So maybe they'll start thinking about secure payoffs. And what we want to give a name to is, what is this best worst-case gain for player 1? So for player 1, let's look at this quantity. It's often helpful-- we'll call it v lower bar for reasons that'll become clear-- is the max over sigma 1 of WG of sigma 1.

So we think of this as the gain, the best gain that player 1 can secure. So if we go back to our definition over here, if player 1 plays a security strategy, then their worst-case gain from that security strategy will precisely be v lower bar. If they chose a non-security strategy, the worst-case gain would be strictly smaller.

But under a security strategy, the worst case gain is exactly equal to v lower bar. So one way of saying-- of defining a security strategy is precisely that the worst-case gain is equal to the maximum possible worst-case gain. And then we can also define, for player 2, what we might call v upper bar.

And that is going to be the minimum over sigma 2 of the worst-case loss of sigma 2. So maybe this is the best worst-case loss or the best loss that P2 can secure.

So what this says is suppose player 2 plays a security strategy. They know that, however their opponent plays, their loss is going to be no worse than v upper bar. And player 1 knows that if they play a security strategy, their gain is going to be no worse than v lower bar. They're always going to get at least v lower bar.

And now we can see it starts to get a little confusing as we look at the min of the max. But what we'd like to understand is the relationship between these two numbers. So the gain that player 1 can secure and the loss that player 2 can secure.

And let's go down here and try to think about this. So the first question we want to understand is, what is the relationship between v lower bar and v upper bar?

Well, the way I've defined them kind of makes it clear. v upper bar is going to be bigger than v lower bar. But let's try to understand why.

What we want to argue is that v lower bar must be smaller than v upper bar? And say again--

AUDIENCE: Is that strict?

IAN BALL: We'll see later. It not be strict, but there's going to be a weak inequality. And one way I like to try to think about this is with a little graph here. Well, let's think of P1 and P2.

And in the end, these will be equal. But let's not draw it this way. What does v lower bar mean for player 1? Let's think about how to interpret it.

It means player one can play a security strategy. And if player 1 plays a security strategy, however the other player plays, their payoff is at least v lower bar. So what my security strategy tells me, this gets to the idea of a guarantee.

It says if I'm player 1 and I play my security strategy, I don't know how my opponent will play. But however they play, my payoff is going to be somewhere in here. It's going to be at least v lower bar.

And the fact that I'm choosing a security strategy means there's no way I can make this graph any better. I couldn't choose another strategy that would ensure that my payoff was strictly higher than this. And similarly, for v upper bar, what it tells me is if player 2 plays their security strategy, it tells me that their loss is always smaller than v upper bar.

So they don't know exactly how much they're going to lose. It depends on what the other player does, but however the other player plays, they're never going to lose more than v upper bar.

Now, this is a bit of a misleading diagram because, in the end, we're going to show these two numbers are equal. But I think this kind of illustrates what goes on.

Now, how can we show this inequality? What we want to understand is, what would happen if each player played their security strategy? So I want to give a graphical argument, and then we'll give the mathematical argument.

But intuitively, suppose we both play our security strategies or both play a security strategy. We may have multiple security strategies. What if both play security strategies?

Well, what's going to happen to player one's payoff? They're playing one of their security strategies. So the gain they get has to be somewhere in here.

But now player 2 is also playing their security strategy, so the loss they experience must be somewhere in here. But they're playing these strategies. There's some actual gain and actual loss.

So what it tells us is we must be somewhere in this box if we both play our security strategies. And let's try to show that mathematically. So let's look at what happens. So here's the result.

For any security strategies sigma 1 star and sigma 2 star, let's try to understand, what is the payoff u1 star u2 star? So player 1's playing their security strategy. Player 2's playing their security strategy, and we want to understand u, which is the gain for player 1 and the loss for player 2.

Well, what do we know about this? We exactly want to go through this reasoning here. Well, the loss, this is the loss for player 2. So this loss can't be any worse than player 2's worst loss from playing sigma 2 star. So this must be less than or equal to the worst loss of sigma 2 star.

Why? Well, this is just the definition. This says if I play sigma 2 star, however my opponent plays, my loss is smaller than the worst loss.

In particular, if my opponent plays sigma 1 star, then my loss must be smaller than my worst-case loss because the worst-case loss is as high as possible. Then, in the other direction, we want to claim this.

Well, if I play sigma 1 star, my worst-case gain is WG of sigma 1 star. So that says, however my opponent plays, I'm going to get a payoff of at least this. So in particular, if my opponent plays sigma 2 star, my payoff is going to have to be higher than this.

And this is v lower bar, and this is v upper bar. Notice here that these middle inequalities didn't actually depend on the fact that these were security strategies. The middle inequalities just followed from the definitions of worst-case loss and worst-case gain. But the outer equalities are using the fact that sigma 1 star is a security strategy for player 1, and sigma 2 star is a security strategy for player 2.

I think there's one last way of interpreting this that may be a bit more reasonable, a bit easier to interpret. And we can think of this as, what would happen if player 1 goes first? And their move is observed? And their mixed strategy is observed?

Well, I claim that if player 1 goes first and plays sigma 1 star, the worst-case gain is exactly what would happen. Why? Well, if I'm player 1, and I go first, and I choose sigma 1 star, and my opponent can see that, well, my opponent and I have exactly misaligned interests. So my opponent is going to try to minimize my gain. And that's exactly what the worst-case gain from sigma 1 star captures.

And if I anticipate that, I want to choose the strategy that will do best, taking into account the fact that, when I do it, my opponent is going to choose the strategy that makes me as worse off as possible. And that's exactly what v lower bar represents.

Similarly, v upper bar, we can think of that as representing, well, if P2 moves first and their mixed strategy is observed, well, now it's the same reasoning. If I observably choose sigma 2, then my opponent, player 1, is going to try to make my loss as player 2 as large as possible. They're exactly going to induce my payoff of the worst-case loss.

Anticipating that, I want to try to make this worst-case loss as small as possible by playing a security strategy, sigma 2 star. So the inequality we've shown here captures a pretty intuitive idea that moving first and having your move observed is a weak disadvantage. Player 1 does worse when they move first, and the other player gets to see what they do, at least weakly, than if player 2 moved first and player 1 got to observe that player 2 was going.

In zero-sum games, this makes sense. Our interests are opposed. If I get to see what you do, that's really going to help me in general. It can weakly help me.

Any questions on this? I feel like this always is quite confusing the first time you see it. Any questions?

Now, I think a lot of people have this intuition that moving first is a strict advantage-- or a strict disadvantage, that if I have to move and the other person gets to see what I'm doing, they can punish me and exploit me, and I'm not going to do very well. And the surprising result that von Neumann proved that-- we're now going to state-- is that it actually doesn't matter who goes first. That these two values, v upper bar and v lower bar, are actually equal. And that's exactly the result we're going to show.

So this is a very famous theorem called the minimax theorem. And this is proven by von Neumann in 1928, so notice Nash's theorem was proven in 1950. So this work on zero-sum games came before Nash. And so Nash, people say, is the founder of game theory, but really, what he did is extended game theory away from zero-sum games to non-zero-sum games. Zero-sum game theory was actually already understood when Nash came along.

And the theorem is, in any finite two-player zero-sum game. Well, v lower bar equals v upper bar. And now let's just write out exactly what that means. So v lower bar is the worst-case gain of sigma 1 star. So this is the max over sigma 1.

Let's make sure this is what we mean to say. Let's look at each of these inner things. If player 1 plays sigma 1 and player 2 chooses sigma 2 to minimize player one's payoff, this is the worst-case gain for player 1 from choosing sigma 1. So this is exactly WG of sigma 1. And indeed, if we then maximize the worst-case gain over sigma 1, this is exactly the definition of v lower bar.

And then, on the other hand end, if I'm player 2, and I'm playing sigma 2, and my opponent chooses sigma 1 to maximize my loss, this is my worst-case loss from playing sigma 2. And if I then minimize this over all sigma 2, that's exactly v upper bar by definition.

Now the inequality from here to here is what we already showed. That's what we showed in class, that this is less than or equal to this. The content of the theorem is to say that, in fact, we have equality. So this graph I showed up here is wrong. In fact, everything just collapses right to here.

And let's go through how we interpret this in terms of who goes first. On the left-hand side, what we imagine is player one moves first. Their move is observed, and then player 2, seeing what player 1 does makes player 1 as worse off as possible.

So this is the case where player one moves observably first. This represents what happens if player 2 moves observably first. And seeing player 2's move, player 1 then chooses what's optimal for player 1.

So what we said is that, in general, we would expect that moving first for player 1 would be a disadvantage. And indeed we showed that this is less than or equal to this. What von Neumann tells us is that if we play optimally, in fact, it's not a disadvantage at all. It's going to do exactly as well as if the other player moved first.

Now, this common number here that's equal to v lower bar and v upper bar, we give it a name. And we'll call this v. It's a good name for something that's equal to both v lower bar and v upper bar. And this is called the value of the game.

So let's try to understand what this means, and let's go through a few examples. So what is the value of a game? We can think of it in a few ways.

What I think it does is it captures-- it's an objective measure of the relative advantage of the two players in this game. So maybe I would say of the fairness or relative advantage in the game.

So we have to put ourselves in the position of before von Neumann's theorem was proven, what people thought is, you had these games, like chess and checkers, and there's not really a way to say who has an advantage. People would say, well, does player 1 have advantage or player 2? Well, it depends who's better at the game. That's what most people would've said. There isn't some single objective number that captures how much of an advantage this game has.

Von Neumann's theorem tells us that, in fact, no, there is an objective measure that tells us, if people play, quote, unquote, "optimally," who should win the game. In other words, how much better player 1 should do than player 2 if both players play optimally. And this captures how fair the game is for the two players or more precisely, quantifies the size of the advantage that either player 1 or player 2 has.

So maybe, at first, you might say, ah, we've just done some inequalities. There's nothing fancy here. But one amazing implication of this is that chess has a value. Chess either-- in fact, we can say more. The value in chess is actually either plus 1, minus 1, or 0.

Now, chess is-- yeah, chess is still a finite game, so everything applies. What this means is if the value of chess is 1, that means that player 1 can guarantee that they win in chess. So who moves first in chess? White? I should know this.

The first player, right? Negative 1 means player 2 can guarantee that they always win in chess, the player with the black pieces. And 0 means each player can guarantee a draw.

And von Neumann's theorem tells us that chess has a value, and in fact, if a bit careful, it tells us that its value must take one of these three numbers. But we don't know which number it is. But which one is unknown?

How is this unknown? I just wrote it down. You just have to do max-min.

Why is this unknown? I mean, I give you-- here's the formula. Just write it down.

The issue is that chess is extremely complex. So the set of strategies in chess is bigger than the number of atoms in the universe. So yeah, we could just try to solve this optimization problem, but it would have an absurdly large number of variables. We wouldn't be able to solve it.

But this, I think, shows us the power of von Neumann's theorem, that it tells us that, in fact, chess has a value. It either advantages one player or the other. We just don't know what that value is because we're not able to calculate it because it's so complicated.

Let's go through a few examples of games to try to understand this concept of a value. So let's start with a game, let's say, top, bottom, left, right. And let's make this 1, 1, negative 4, negative 8.

So player 1 chooses top or bottom. Player 2 chooses left or right. And the question is, what is the value in this game?

Or another way of saying this is if you were going to play this game and you were asked, do you want to be player 1 or player 2? Which player would you rather be? Any thoughts? Yeah.

AUDIENCE: I forgot which-- which is 1 and 2?

IAN BALL: This is player 1, and this is player 2.

AUDIENCE: I'd probably prefer to be 1.

IAN BALL: OK, and why?

AUDIENCE: [? You ?] [? can ?] work it all out [INAUDIBLE]. I don't [INAUDIBLE], but if you did the expected value [INAUDIBLE].

IAN BALL: But I don't know, negative 4 or negative 8, these are pretty big negative numbers. So if I just look at it, I see some really big negative numbers. You're right, but I'm not convinced. You haven't convinced me yet.

But let's notice if we just naively looked at this game, we might say, oh, this is a good game for player 2 because, remember, a big negative number for player 1 means player 2's doing really, really well. But the structure of the game affects who has an advantage.

Anyone else, how would they play. Maybe in the red sweatshirt? I don't think I've heard from you. Yeah.

AUDIENCE: Player one plays [INAUDIBLE].

IAN BALL: Exactly, so even though we have some really big negative numbers down here-- and notice this is still a zero-sum game. I think people get confused, and they say, oh, we have these really big negative numbers. It's still zero-sum. What this negative 4 means is player 2 wins 4 if we come down here.

But because the structure of the game, if I'm player 1, my security strategy here is T. If I play T, whatever my opponent does, I'm going to get a payoff of 1. And if I were to play anything else other than T, I wouldn't be able to secure a payoff of 1. And in this case, it turns out player 2, actually, anything is a security strategy for player 2.

And so what we see in this game is the value v is 1. Player 1 by playing T can secure a payoff of 1, and there's no way they can secure anything higher than that. So hopefully, this makes it a bit more concrete what these numbers mean. Let's do another example with a similar flavor.

Say, well, that's a bad example. Well, it doesn't really matter. We'll make all the numbers positive. That's fine. Let's look at this game.

What is the value of this game? Or which player-- and I guess, here, it's pretty clear which player you'd rather be because all the payoffs are positive. So I definitely want to be player 1 here, but the question is, what is the value of this game? Think through it, yeah.

AUDIENCE: It'd be negative 8.

IAN BALL: Negative 8? OK, I don't think so, but let's think through it. So why negative 8?

AUDIENCE: The same player 2, [INAUDIBLE] you could guarantee the results [INAUDIBLE], positive.

IAN BALL: This is where it gets really confusing. So player 2 is trying to-- they want to minimize things. So let's think through. So let's just walk through it.

Let's say player 2 plays R. Then they're either going to lose 23 or lose 8. So they would want to choose L, and if they choose L, well, whatever the other player does, the loss to player 2 is only 1.

So playing L is much better. So it turns out, in this case, L is going to be a security strategy for player 2. And the value of the game is then going to be 1. Player 1 gets 1, and player 2 loses 1. So here, we see the value is 1.

Now, notice an implication of this. If we looked at these two games, if you just, at first, stared at it, you would probably say, this game is so much better for player 1 than this game is. Because if I look at every cell of the game, the payoff is at least as high in this game.

And sometimes a lot higher. 23 is a lot bigger-- or 8 is a lot bigger than negative 8. 23's a lot bigger than 1. But von Neumann's theorem actually tells us, no, each of these games has some objective value that captures the relative advantage of player 1 and player 2 if people play optimally. And in fact, these two games have exactly the same value even though a lot of the payoffs seem a lot higher in this game.

Now, would you really be indifferent between these games? Well, it depends who you're playing against. If you're playing against a computer who's really rational, probably. If you're playing against a person who makes mistakes, maybe you'd actually rather play this game because, when we calculate the value of the game, there's this implicit assumption that people are behaving-- are very good at computing things and are playing rationally. OK, a final example.

What about this game? What is the value of this game? And what would the security strategies be? Yeah. It is, yeah. And what would the security strategies be here?

AUDIENCE: [INAUDIBLE]

IAN BALL: Exactly, so let's go through it. So here, the value is 0. This is a game we've studied before. We've just only put in one number here.

Let's say I'm player 1. Let's try to compute my security strategy. Well, if I play T, things could go really badly because my opponent could play R, and I get negative 1.

If I play B, things could go really badly. My opponent could play L, and I would get negative 1. But if I mix 50/50, then I ensure that I always get a payoff of at least 0, in fact, exactly 0.

And we can do similar reasoning for player 2 and see that the value is 0. And in fact, each player has unique security strategy. So security strategy equals 1/2, 1/2 for each player.

I guess, technically, it's 1/2 T, 1/2 B for player 1 and 1/2 L plus 1/2 R for player 2. So I'm being a little vague here. Great. Yes.

AUDIENCE: Is it too general to say that there's not a strictly dominant strategy that [INAUDIBLE] always going to be 0?

IAN BALL: No, because-- no, that's not going to be true. So actually, in these cases, we only have weak dominance, not strict dominance. But in general, you can have much more complicated games where-- so we'll actually see an example later today of a game with that property. But good observation.

In fact, one player could even have a strictly dominant strategy and still lose a lot. It be that one strategy's better than everything else. But whatever I do, I lose money. Yeah, any other questions?

OK, so now we would like to connect what we've done here to what we've been doing throughout the course. So I think it's important to understand that, before, what we looked at was Nash equilibrium. And Nash equilibrium captured stability. It said that if each player is playing according to a Nash equilibrium, then this strategy profile is stable in the sense that no player can do strictly better by unilaterally deviating.

Security strategies capture something very different. Security strategies, well, they capture security or safety. Each player is just thinking, how can I play to ensure that, whatever happens, I do reasonably well?

But it turns out that there's a tight connection between these, and that's what our next result is going to say. It's going to say that, in fact, in zero-sum games, these are basically the same concept, that Nash equilibria and security strategies coincide in zero-sum games. Let me be more precise about that.

So in, as usual, any finite two-player zero-sum game-- now, remember, Nash equilibrium always makes sense. Security strategies only make sense in two-player zero-sum games. So it's not even meaningful to say, do these concepts coincide beyond zero-sum games? Because this concept doesn't even make sense.

But in any finite two-player zero-sum game, at least both of these concepts make sense. And let's say, for any strategy profile, mixed-strategy profile, sigma 1, sigma 2. So I fixed a game, and I consider an arbitrary strategy profile, sigma 1, sigma 2.

I'm going to argue that sigma 1, sigma 2 is a Nash equilibrium if and only if-- so I'm saying two statements are equivalent. One statement is that this strategy profile is a Nash equilibrium. The other statement is that each component of this profile is a security strategy. So if and only if sigma 1 is a security strategy for player 1 and sigma 2 is a security strategy for player 2.

So this is an "if and only if" statement. So let's make sure we understand both directions. It says, if you give me a strategy profile in a zero-sum game and it is a Nash equilibrium, I can immediately conclude that player 1's strategy must be a security strategy for player 1, and player 2's strategy must be a security strategy for player 2.

Now let's look at the other direction. It says if I find a security strategy sigma 1 for player 1, and I find a security strategy sigma 2 for player 2, and I now consider the strategy profile containing these two strategies, so player 1 plays his security strategy, player 2 plays her security strategy, then that strategy profile is necessarily a Nash equilibrium.

So it connects these two ideas. It also has an implication for computation. Let's say you're asked on a problem set in an exam to find a Nash equilibrium of a zero-sum game. This theorem tells you one way you could do it. How could you do it?

AUDIENCE: I think you'd find a [INAUDIBLE] strategy.

IAN BALL: That's it. You go to player 1. You solve for player 1's security strategy. There might be more than one, but you find one of them.

You go to player 2. You solve for player 2's security strategies. You find one of them. You put those together. Now you have a Nash equilibrium.

Now the nice story is that this is computationally easier than finding Nash equilibrium the other way. That's true for computers. I find this actually very rarely true on exams, that you're better off taking this approach, but if you were a computer, you would find that this would be a very good way of computing Nash equilibria. But for humans, I don't know. It's hard to find a good example.

But in principle, this gives easy computation of Nash equilibrium. And what's amazing about it is that you can compute the Nash equilibria by separately finding a strategy for the two players. Normally, Nash equilibrium is about the interaction between strategies.

And what makes finding it hard is what's best for one player depends on what's best for what the other player does. But this says you can separate this computation into just two simple optimization problems for the two players. And then you get a Nash equilibrium. So computationally, that's why it's very nice.

OK, I want to try to show you the argument for this. It's very short, but it is-- have to be a little careful. So let's go through it.

So let's first show this direction. So what I want to say is suppose I've computed a security strategy for each player. So let's say sigma 1 star is a security strategy for P1, and sigma 2 star is a security strategy for P2. I've done that.

Now what I want to argue is that this pair, sigma 1 star, sigma 2 star is a Nash equilibrium of the game. And really, I just have to remember my formula over here. So let me just look at my formula from over there. I want to look at u of sigma 1 star, sigma 2 star.

And I immediately know that this is-- should I do it the other way? Which way are we doing? It might help. It might look nicer if I do it the other way.

We'll start here. OK, now, up here, we just had an inequality between these statements. But now we have an equality because of the minimax theorem. So the minimax theorem tells us that the worst-case loss from sigma 2 star must equal the worst-case gain from sigma 1 star. And we have equality in this expression from that last inequality.

And now what we want to show is that we have a Nash equilibrium. Well, let's look at what happens if player 2 were to deviate. So this is what happens under our strategy profile. Let's suppose that player 2 deviates to some strategy sigma 2 prime. How does this relate to the worst-case gain for player 1? What can we say about the relationship between this and this? Yeah.

AUDIENCE: [INAUDIBLE] gain is at least [INAUDIBLE].

IAN BALL: I think it's actually going to be this way. Let's make sure. So this is always really confusing. The worst-case gain says, whatever my opponent does, I'm guaranteed to get a payoff of at least the worst-case gain.

So one thing my opponent could do is they could deviate to sigma 2 prime. And if they chose sigma 2 prime, then that's one thing they can do. It must be at least as good as the worst thing they can do for me as player 1.

OK, now, let's try to get this last one. What if player 1 were to deviate? So player 1 deviates. What's the relationship between this and the worst-case loss for sigma 2 star? It's really easy to get confused here.

AUDIENCE: Would it be the same inequality [INAUDIBLE]?

IAN BALL: It's actually going to be the same, right? Because this says my-- now I'm taking the perspective of player 2. If I play sigma 2 star as player 2, the worst loss for me is this.

So if player 1 plays sigma 1 prime, that loss can't be as bad as the worst loss, which means the loss is smaller. Not as bad loss is a smaller loss. So here, I've just used the definitions of worst-case loss, and here, I'm actually using the min-max theorem. Because, in general, we only have inequalities here, but the min-max theorem says they have to be equalities.

And now I claim that we're done. Someone look, now we just have to squint. Of course, this proof is pretty hard to come up with.

I've done it just right, so it seems easy. It's not easy. You have to know exactly where you're going.

But can someone just walk me through why I've already shown that this is a Nash equilibrium? If we just look really carefully. I always find this confusing. Yeah.

AUDIENCE: I mean, if you look, that is player 2 doing worse. So they don't want to--

IAN BALL: OK, so are you the left side or the right side? Which side are you on? You're on the left side. So you're here. OK.

AUDIENCE: This is player 2 deviating and then doing-- No, this is player 1--

IAN BALL: This is player 1 deviates. So let's make sure. If we compare here, player 2 is doing the same thing. And player 1, instead of doing what they're supposed to do in equilibrium, is deviating to sigma 1 prime. So player 1's deviating and what happens?

AUDIENCE: They're doing worse [INAUDIBLE].

IAN BALL: They're doing weakly worse. So this says any deviation by player 1 is unprofitable. Player 1 does weakly worse by deviating from sigma 1 star to sigma 1 prime. And crucially, this is holding for all sigma 1 prime and sigma 2 prime.

OK, so what we've shown-- what you've argued verbally is that player 1 cannot profitably deviate. However they deviate, player 1 does weakly worse. What about the other case?

What about player 2? How does this show me that player 2 can't profit? Yeah.

AUDIENCE: So on the right where player 2 is deviating--

IAN BALL: Yes, exactly.

AUDIENCE: And player 1 will be doing at least as bad or better [INAUDIBLE], don't have any alternatives.

IAN BALL: So you have to be careful, though, because we're interested in how player 2 is doing. So player 2 was doing this. They deviate to this.

It looks like player-- that the payoff is higher. But we have to think it's player 2. So what does this mean?

AUDIENCE: 1's payoff is higher, and player 2's is lower, or just--

IAN BALL: Exactly, right. So this says if player 2 deviates from sigma 2 star to sigma 2 prime, this is player 2's loss. So this is player 2's loss is weakly higher than their loss in equilibrium.

That means their deviation makes them weakly worse off. And therefore, neither player can benefit by deviating. And therefore, we have a Nash equilibrium.

So you just have to string the inequalities together in just the right way. I always think it's very easy to get confused, but here we are. Any questions on this?

So now we want to show the reverse direction. We want to show let's start by saying suppose sigma 1, sigma 2 is a Nash equilibrium. We want to say that both players must be playing security strategies.

But we're actually going to-- erase this. We're actually going to prove this by contradiction. So we have to be a little careful.

What we're going to say is to show that if it's an equilibrium, both players must be playing security strategies. We're going to show if some player were not playing a security strategy, then it could not be an equilibrium. So by symmetry, let's say suppose sigma 1 is not a security strategy for player 1.

I Want To Show, WTS, Want To Show that sigma 1, sigma 2 is not a Nash equilibrium. So this is my way of saying, in any Nash equilibrium, both players must be playing security strategies because if some player is not playing a security strategy, then it couldn't have been Nash equilibrium.

Now, I'm saying, what if sigma 1 is not a security strategy? There's a symmetric argument if sigma 2 is not a security strategy. So one way of saying this is I can't find an equilibrium where players are not playing security strategies. That's one way of saying this.

And now this one-- again, it's short. The argument's really short if you know exactly where you're going. But it's tricky. I need to make sure I get it right in my notes.

So what do we know? Let's look at sigma 1, sigma 2. This is always what we look at. And what's the relationship between this and WG of sigma 1?

AUDIENCE: Is that [INAUDIBLE]?

IAN BALL: I think it's at least. Let's see. Maybe that's what you said. Maybe I didn't hear you. Let's check.

This is the-- if I play sigma 1 as player 1, this is the worst that can happen to me as player 1. So sigma 2 is one thing player 2 could do. The one thing they could do is at least as good for me as the worst thing that could happen, so I get this inequality right.

And because we have this weak inequality, we're going to argue in two separate cases. There's two possibilities. Either this is strictly bigger than this, or it's equal OK. So let me actually write it down like this.

Write it like this. So what I'm saying is there's two cases. Maybe you can add a bit more detail on your notes to make it clear, since I'm of doing it dynamically. But we know that this is weakly bigger than this.

So there's only two possibilities. Either they're equal, or one is strictly bigger than the other. And I want to show that, in either case, we don't have a Nash equilibrium. In either case, someone has a profitable deviation.

I think this one is actually the easier direction. So here, I want to show that player 1 has a profitable deviation, using the fact that sigma 1 is not a security strategy for player 1. Can anyone try to think through-- I think it's good to try to think through it before you see it. Otherwise, there's no way to remember it.

So I want to claim-- I want to argue that player 1 has a strictly profitable deviation, precisely because sigma 1 is not a security strategy. Anyone think through how we might do that? Yeah.

AUDIENCE: Perhaps the worst gain of sigma 1's security strategy is going to be greater than the worst case or the worst gain of this current strategy that sigma 1's playing, which means that if we switch to the security strategy, we're getting a strictly-- or weakly, I guess, [INAUDIBLE].

IAN BALL: It's actually strictly, yeah, so great. So exactly right. Sigma 1 is not a security strategy.

So let's consider a deviation to a security strategy. What happens if player 1 deviated to a security strategy, say, sigma 1 star? OK.

So I want to say, instead of playing sigma 1, which we know is not a security strategy, let's say player 1 deviates to a security strategy, sigma 1 star. And I want to show that player 1 does strictly better. How do I show it? Well, because sigma 1 is not a security strategy, the worst gain from sigma 1 is strictly worse than the worst gain from the security strategy.

This is what it means to be a security strategy. It means sigma 1 star maximizes the worst gain whereas sigma 1 does not. Now I'm almost done. Now I just need one more inequality here. How do we get this last inequality?

This is the same thing we've always been using. But it's always hard to remember. If this is the worst that player 1 does when they play sigma 1 star, then if player 2 plays sigma, player 1 must do weakly better than the worst thing.

And now if I tie these inequalities together, what have I shown? I've shown that player 1 can strictly benefit by playing sigma 1 star. I only have a weak inequality here, but I have a strict inequality here. So if I link them all together, I've identified a strictly profitable deviation for player 1. I'm saying if player 1 is not playing a security strategy, they must be able to profitably deviate to a security strategy.

Though, this relied on this assumption here. Now let's go to the other case. So here what I showed is that player 1 has a profitable deviation. So any guesses about what I might show here? Yeah.

AUDIENCE: Player 2 has a profit.

IAN BALL: Player 2 has a profitable deviation. So I want to show that player 2 has a profitable deviation. Hmm, any thoughts?

How can I-- well, we actually just have to unpeel the definition of the worst-case gain. What is this? Well, this is exactly the minimum over sigma 2 of u of sigma 1, sigma 2.

And I claim we're already done. We just have to look at it really carefully. Ah, sorry, sigma 2 prime. How can we see that player 2 has a strictly profitable deviation? Yeah.

AUDIENCE: Considering that currently, the sigma 2 doesn't actually allow the utility to be minimized, there's some other sigma 2 prime that actually achieves the minimum that sigma 2 can be replaced with.

IAN BALL: Exactly, right. Let's choose a sigma 2 prime that achieves this minimum. Then the value of this is going to be strictly smaller than the value of this. But remember, we're talking about player 2.

So that means player 2 has a deviation to sigma 2 prime that makes their loss strictly smaller than their loss here. And that means player 2 profits.

OK, let me stop there. This is one of these proofs where I think you read it in a textbook, and it's like two lines. But unless you go over it really, really slowly and carefully yourself, you're never going to understand it.

AUDIENCE: Just to clarify, are you essentially saying that then W2 will switch to the case where you're in [INAUDIBLE]-- in the first case, player 2 will switch to whatever strategy [INAUDIBLE] you're in the second case. And then you have [INAUDIBLE]?

IAN BALL: No, it's not the case that we're going to immediately get an equilibrium if player 2 deviates. What's tricky is that, here, it could be that, because player 1 is not playing a security strategy, player 1 might be doing something really weird here. And it could be that, given this really weird play by player 1, the best thing for player 2 is actually something else really weird. So it doesn't necessarily mean that, when player 2 deviates, we're going to get an equilibrium.

This is not necessarily a security strategy for player 2 because it's the strategy that does best against sigma 1, which could be something really, really strange, and therefore it's not necessarily a security situation.

AUDIENCE: OK, [INAUDIBLE].

IAN BALL: OK, good, yes.

AUDIENCE: Is there any way to [INAUDIBLE] this?

IAN BALL: Is it pretty clean? I don't know, what do you mean by clean? Yeah.

AUDIENCE: I'm just like, because this is proof by contradiction.

IAN BALL: Oh, I see. Yeah, I mean, I think I went for the shortest possible proof. I mean, whenever you have a proof by contradiction, you can always undo it.

I mean, here, everything's finite. So there's not-- there's these math philosophy things about proofs by contradiction. But here, I think you can just pretty easily convert it to a, quote, unquote, "direct proof." We can talk more about that.

OK, so this has been kind of math heavy. Let me end the last four minutes with a maybe more fun application of this. So in the past, I went through a really simple game that we could solve exactly. Here, I want to go through a more realistic game and just show you what the answer is because I think it's a nice demonstration of the power of the theory.

So poker is too hard to solve, full poker. So we're going to look at a really simplified version of poker that von Neumann actually came up with that's called von Neumann poker.

And here's how it works. Player 1 is going to get a hand. So P1 gets a hand, which is just a number x between 0 and 1, where higher numbers are good. So instead of thinking of an actual deck of cards, player 1 is dealt some hand-- a number between 0 and 1, and it states uniformly distributed.

Player 2 is dealt a hand, Y, That's also in 0, 1. And then they're each going to ante $1. So each player puts $1 in the pot. And then we're going to play a very simplified version of poker where player 1 chooses whether to bet or check.

Checking means not betting. Betting means betting. And there's only one amount they can bet, so the bet size is going to be B. And that's some fixed number.

It could be 1. It could be 2. It's just a fixed number.

Now if player 1 bets, then player 2 can either call or fold. Really simple version of poker. And then what happens? Well, if player 1 checks, then whoever has the higher hand wins the pot.

So I'll just write, it's a little vague, plus 1 minus 1. And this means if I check, we just compare our hands, and the higher hand gets one, and the lower hand loses one. If player 1 bets and player 2 folds, well, player 1 automatically wins. So player 1 gets a payoff of 1.

And then, down here, if player 1 bets and player two calls, now each person has put, one, their ante plus their bet B into the pot, and the person with the higher hand wins that, and the person with the lower hand loses that. So I'll write down here plus 1, b plus 1. So when I write plus minus 1, it means the payoffs depend on who has the higher hand. If player 1 has the higher hand, they win their ante-- or really they win the other player's ante and the other player's bet. If they have the lower hand, then they lose their own ante and their own bet. So they get minus.

And the question is, how should you play? And who has an advantage in this game? Here, I'm assuming risk neutrality.

So let's say the value of the game and how to play. So when I ask the value of the game-- let me start with a more informal question. It's who has the advantage?

And I can formalize that by asking, is the value of this game positive or negative? So who has the advantage is the vague question. The mathematically precise question is, what is the value of this zero-sum game?

How to play. This is a vague question. The mathematically precise thing is, what are the security strategies for the two players? So

You can see how the theory presented today gives a formalization of these very-- of these vague concepts. And we're running out of time, so maybe I'll put this on the problem set to solve it. But just maybe final thoughts.

Would you rather be player 1 or player 2 in this game? I think it's not at all clear. Any thoughts? Yeah.

AUDIENCE: I'd rather be player 1.

IAN BALL: Why?

AUDIENCE: Because then I get to set the rules of the game [INAUDIBLE].

IAN BALL: Yeah, so here the bet is exogenous. You can't control the size of the bet, but you do control whether you bet. And you're actually right. In this case, it is going to be player 1 who has an advantage. But the reason is actually this strange thing here that if player 1 checks, player 2 doesn't have the right to bet.

In normal poker, player 2 can then bet after a check. And in standard poker, you actually have an advantage generally going second and going later. Here, player 1 actually has an advantage.

So it turns out the value is greater than 0. And how to play, I'll just say the punchline is that, it turns out, player 1 sometimes bluffs. Meaning player one sometimes bets even when their hand is really, really bad.

And I think, before this, people thought poker is this psychological game. Bluffing is this thing that's beyond mathematical reasoning. And it turns out just the minimax theorem immediately explains why bluffing is sometimes a good thing to do. But I'll put this on the problem set for you to work out.