[SQUEAKING]

[RUSTLING]

[CLICKING]

**ANKUR MOITRA:** So let me tell you about zero sum games. Let me start off with the key definition. And we're going to study what happens under the simple definition. So first, let me tell you the abstraction. And then we'll do some examples to try and make it concrete. So specifically, we'll be interested in what's called a two-player zero sum game.

And a two-player zero sum game has a few components. So it's described by the following things. First of all, there is a row player. A lot of times, we'll call this player Alice. And Alice has m different strategies. These are the moves that she's allowed to make in the game we're playing.

There's a second player, which is called the column player. And usually, we call the column player Bob. These are just the standard conventions. And Bob is going to have n strategies. So already, you can see that the game looks different for the two players. They have even different action spaces of what they're allowed to make as a move in the game.

And the crucial thing is the last component, which tells us what are the outcomes of the game-- who pays who how much depending on what strategies you play. So finally, what we have is we have an m by n matrix, which is called the payoff matrix. And we'll call this C. And the way that this matrix C works is that the entry Cij is the amount that the row player, i.e. Alice, gets in payment from the column player if they play i and j respectively.

So notice that the row player, Alice, has m strategies and there are m rows in this matrix. The column player, Bob, has n strategies. And there are n columns in this matrix. So you should think of this payoff matrix as being indexed by their strategy set. Every choice of a strategy for the row player just selects which row we're interested in. Every strategy for the column player just selects which column we're interested in.

And then we look into that entry in the matrix. And that describes the net payment from Bob to Alice. That payment could be negative, in which case the money goes from Alice to Bob. But in any case, right now, in the case of a zero sum game, we only have a single payoff matrix that describes what happens. And it's called the zero sum game because, if you add up the wins and losses of Alice and Bob, it's 0.

So of course, you can consider more general games. And we'll talk more about that towards the end. But this will be our basic abstraction. So let's see some examples to get some intuition for this. There's a very simple example that can get us started. This is called the matching pennies game.

So the way that this works is that Alice and Bob simultaneously choose either H or T. So in particular, their strategy space-- m equals n equals 2. They only have two choices of actions. They can choose either heads or tails. And now, what we have to do is describe what the payoffs are.

So if they match, if these two coins match, then Alice gets $1 from Bob. And otherwise, if they don't match, then Alice pays Bob $1 instead. So this is a very simple example of a zero sum game. In fact, we can write it in this formalization that I described right here just by writing down what the payoff matrix is. So in this case, the payoff matrix is just the 2 by 2 matrix.

And in particular, It'll look like this-- 1, minus 1, minus 1, 1, where we have some conventions of what the different rows map to in terms of the strategies. So the first row maps to the choice of Alice choosing heads. The second row maps the choice of Alice choosing tails, and then similarly for Bob, but for the columns.

And so the entries in this matrix just describe the payoffs of this game. So let's make sure we have the right intuition for this. So let's think about the zero sum game. Which player has an advantage, if any? So who would you want to be, if you got to choose to either be the row or the column player? Or would it not really matter? Just asking for your intuition. Yeah?

**STUDENT:** It doesn't matter.

**ANKUR MOITRA:** Doesn't matter? Why?

**STUDENT:** I mean, it's the same payoff for [INAUDIBLE].

**ANKUR MOITRA:** Yeah, that's right. So this is very intuitive, that neither of these players has an advantage. It doesn't matter if you're tasked with playing the role of Alice or Bob. But for more complicated games, it's maybe not so obvious who has an advantage. And this is where we're going to use ideas from linear programming to try and understand who actually has an advantage and in what sense do they have an advantage.

So maybe for simple games, it's easy to reason about them just by staring at them and figuring out who you'd rather be. But for more complicated games, maybe it's not so simple. So let's do a more interesting example. And we'll come back to this example later on, once we've developed the theory.

So let me just write down the payoff matrix. So let's say our payoff matrix C is a 3 by 3 matrix. So each Alice and Bob have three different strategies. But now, the game won't be symmetric in such an obvious way. So the payoffs are going to look like this.

So Alice can either choose the strategy, let's say, a1, a2, or a3. Bob can choose the strategy b1, b2, and b3. And as before, the entries in this matrix describe the payoff of one player to another. So let me get some volunteers. So who wants to play this game? I need two volunteers. Yeah, come on down. Anyone else? No one else wants to be on camera? I'm going to select someone. Don't be shy.

You guys brought some money with you, right? No? OK, that's all right. I'll spot you. I have these convenient baggies of monopoly money. So what I want you to think about is, in this game, who do you think has the advantage? So the row or the column player? And can we make predictions about, if they play rationally, what their strategies ought to converge to? I think this isn't so obvious. But let me start with you. So who do you want to be, Alice or bob?

**STUDENT:** I want to be Alice.

**ANKUR MOITRA:** You want to be Alice? All right. I'm afraid you're stuck with Bob. All right, so what I'm going to have you guys do is grab some chalk. And I'm going to ask you guys to write down which strategy you pick-- a1, a2, a3, b1, b2, b3. Hide your answers from the other person, so they don't see it. And then we'll reveal them simultaneously and figure out who's doing better. So we'll play a few rounds of the game. So pick a spot on the board and figure out, based on the payoff matrix, what exactly you want to play. And we'll see how we do.

**STUDENT:** [INAUDIBLE]

**ANKUR MOITRA:** Yeah, why don't you use this one, so that you guys-- yeah.

**STUDENT:** So I just write--

**ANKUR MOITRA:** Yeah, yeah, yeah. You can ask the audience for help, too, if you want.

**STUDENT:** Oh.

**STUDENT:** [INAUDIBLE]

**ANKUR MOITRA:** Yeah, maybe cover it up. Yeah. OK. Are you guys ready? All right. Reveal your answers. b1 and a1. All right. So, Alice, congratulations. You win $2. Can you pay the lady?

**STUDENT:** Thank you.

**ANKUR MOITRA:** All right, let's do another round, since maybe you guys are just getting the hang of it. All right, so think hard. Last time, she played a1. So maybe you want to change your strategy. I don't know. That's up to you. a1 worked out pretty well. I don't know if you want to change. So think hard, and then write down your answer.

So we've got a1 and b2. So you owe him $1 now. He's learning. You're still up. But now, some of your winnings are going away. Let's do another round. What do you guys want to play this time?

OK All right, so let's see. b3 and a2. Ooh, you lose three more dollars. Uh oh.

**STUDENT:** How much do I have?

**ANKUR MOITRA:** You have $7 total. You're going to go bust soon, though, if you keep it up. All right, let's do another round. Let's give her a chance to get back on top. All right. So write down your next strategy. OK, b3 and a1. All right. So oh, she just won it all back.

All right. So you guys are pretty close to even. So let me just test your guys' intuition. You're not necessarily supposed to have intuition about this yet. So after seeing them play, who thinks Alice has the better deal? You have to vote for something. It's OK if you're wrong. We'll build up intuition as we're going along.

Who thinks Bob has the better deal? All right, let's try that again. I got a grand total of like two votes. Who thinks Alice has the better deal? All right. Who thinks Bob has the better deal? All right, it's a pretty mixed. All right, so thanks for your help. And we can figure out the theory of how zero sum games work. And we're going to be able to answer this question.

But for things like matching pennies, as I told you, it's very intuitive. You guys all got the right answer about who has an advantage. They're both the same. But as soon as you move to bigger games-- and this isn't even much bigger. This is just a 3 by 3 game. Already, it's not so obvious who has the advantage.

So one of the things that we're going to do at the very end of lecture is we're going to talk about really gigantic games, games which I could never even write down on the board, because there are all kinds of games that show up in things like advertising for presidential elections.

Maybe you want to figure out how to allocate your total advertising budget across a whole bunch of battleground states to figure out how to win the most electoral votes. These are exactly the kinds of things that can and are modeled as gigantic zero sum games. And we're going to develop the theory for how to think about these things. So it's good that we don't have intuition because we can make progress.

All right. So let me introduce a little bit more terminology. Let me tell you what it means to solve a game. And there will actually be a few different notions of how to solve games. And they'll all behave a little bit differently. So the easiest concept to wrap your head around is called a pure Nash equilibrium.

So this is a concept that makes sense regardless of whether the game is zero sum or not. For general two-player games, I could have two payoff matrices, one of which that describes what Alice wins and one of which describes what Bob wins. But it's not necessarily a transfer of money from one to the other. So a lot of the things we're doing, though, the only way to think about it as a linear program will be if we're working in the zero sum game, even though this concept makes sense more generally.

So what is a Nash equilibrium? It's a pair of strategies i, j-- so this means that the row player plays i. The column player plays j. And it has the following property, which I'll just describe informally. But then we'll make it mathematically precise. It has the property that no player has any incentive to deviate.

How many people have seen the movie*A Beautiful Mind?* It butchered this concept. So please do not get that definition from that movie. So we'll do the correct definition here. So what do I mean by incentive to deviate? Let me turn this into a set of linear inequalities. That way, we can think about it in a more mathematically precise way.

What I claim is that the row player, if he's playing strategy i, and the column player is playing strategy j, then the row player has no incentive to deviate if and only if, for all i prime, Ci prime j is at most Cij. So the way to think about this inequality is that the column player has already committed to j. That's what we mean by i, j being a pure Nash equilibrium, is that row player's decided I'm going to play i. He's even announced it. I'm playing i.

The column player has announced I'm playing j2. And even with that knowledge of what the other player is doing, they still can't do strictly better than what their current strategy is. So if the column player has announced that they're playing j, the only thing that's relevant is the j-th column of our payoff matrix.

This would be the payoff if the row player sticks to his current strategy. He still plays the strategy he announced, strategy i. And hypothetically, if instead he changed to another strategy, i prime, the claim is that he doesn't get any more money, any positive win for making that deviation. So he might tie, in terms of changing to another strategy. But he doesn't do strictly better.

And similarly, the column player has no incentive to deviate if and only if for all j prime Cij prime is at least Cij. Now, just to make sure people are paying attention, so here, obviously, the deviation is happening over the choice of i versus i prime. Here, the deviation is happening over the choice of j prime versus j. Why are my inequalities flipped here? Why is my notion of no incentive to deviate a less than or equal to constraint here when I change i to i prime? And why is it greater than or equal to constraint when I change j to j prime? Yeah?

**STUDENT:** The payoff is like [INAUDIBLE] player, so it's the [INAUDIBLE]

**ANKUR MOITRA:** Exactly. That's right. So really, the payoff for the column player is just negative of this. So when I multiply through by minus 1, everything switches. So this is a very natural notion. This is called a pure Nash equilibrium. But already, we can see some deficiencies of pure Nash equilibrium. So even if we come back to this matching pennies game, so is there a pure Nash equilibrium? What do you guys think? I see someone shaking their head in the correct direction No, there isn't. Can you tell us why? Any intuition?

**STUDENT:** I guess if if I'm the rogue player, I always want to choose one that will give me 1. And then if the tails person switches their strategy, I'll switch to the other one.

**ANKUR MOITRA:** That's right, that's right. That's exactly right. Let me just rephrase it slightly. A pure Nash equilibrium corresponds to announcing what your strategies are publicly and still there being no reason to switch. So whatever those strategies are, imagine that Alice announces she's going to play heads, Bob announces he's going to play heads.

So then who has the incentive to deviate? Bob. He doesn't want the pennies to match. So if he deviates to tails, then he's going to win $1. And he'll do strictly better. And that's true for any particular starting configuration. There's always one player who has an incentive to deviate. So unfortunately, there's no pure Nash equilibrium.

Now, even though there wasn't a pure Nash equilibrium, we can already intuitively see the way the game would be played. Maybe Alice and Bob just choose randomly. And that would be pretty good because in expectation they would win no money. And that's probably the best that they can do.

But for this type of game, you might have the intuition that I cooked this up in such a way that, again, there's no pure Nash equilibrium. So when there is a pure Nash, it becomes very obvious how to play, for the most part. But when there's no pure Nash equilibrium, all of a sudden, you have to ask how you should play.

And the key point is that maybe the notion of playing shouldn't necessarily be a deterministic choice because the problem is that, when you make deterministic choices, you can't commit in a way where you never have an incentive to deviate. So this takes us to the next notion, which will be critical for us, which is a finer grained notion of an equilibrium, which is called a mixed Nash equilibrium. Basically, it's one which allows randomization. So let me define it. This will be the key thing that we talk about today.

So a mixed Nash equilibrium is a pair. But now, this pair is going to mean something different. So it's a pair x, y with-- and I'll tell you what this means in a second-- x belonging to delta m-- I'll tell you what that means-- y belonging to delta n again with the property that no player has an incentive to deviate, has a positive incentive to deviate.

So I have to tell you what this notation delta sub m means. And what this means is just a distribution on the set m. So the way to think about it equivalently is that x is a vector that has non-negative entries. And the sum of all of these xi's from i equals 1 to m is equal to 1. That's what the constraint is.

So what's going on is that, instead of the row player choosing one fixed strategy, like strategy i, he's going to randomize over different strategies. And the probability that he puts on each different strategy is given by the associated coordinate of this corresponding vector x.

So when I write x belonging to delta m, keep in mind that this is really a set of linear inequalities, which is great because we've been talking about linear programs. It's just the constraint that x is an m dimensional vector, all of its coordinates are non-negative, and they sum to 1.

Same thing is happening here, except with the fact that it's a distribution on n, not m, because the column player is randomizing over the choice of which of the n columns, because that's the number of strategies. So this is the same basic idea as a pure Nash equilibrium. We still have this notion that no player has an incentive to deviate.

But now, we're augmenting the set of strategies. So instead of me announcing I'm playing strategy i and you announcing you're playing strategy j, I'll instead announce that I'm going to sample my strategies from the distribution x. And you will announce that you're going to sample your strategies from the distribution y. And we can ask, even after you make that announcement, if you have any incentive to change up your strategies.

So let's be a bit more precise about what exactly I mean by no incentive to deviate, because this will come up later when we do the proof of a very powerful theorem I'm going to cover very soon. So what does it mean for the row player to have no incentive to deviate now, when we're talking about randomized strategies?

Well, let's first do the following. So first of all, the expected payoff to the row player is the sum over j equals 1 to m-- the sum-- oh, sorry, i, i equals 1 to m, j equals 1 to n of Cij xi yj. And in matrix vector notation, this is just the quadratic form of x and y on this payoff matrix C.

So all I'm doing is I'm saying, all right, you've announced that you're playing strategy x and you're playing strategy y. The first thing I could care about, once you've announced that those are your strategies, is how much does Alice win in expectation. And this is a very simple probabilistic computation because all this is is it's the probability that Alice plays strategy i, because that's given by the i-th coordinate of x, times the probability that Bob plays strategy j.

They sample their strategies from their distribution independently. And then this is whatever the payoff would be. So this is just computing the expected value, the expected winnings that Alice has. And in our case, it works out to a very nice expression in terms of vectors and matrices.

But now, we can talk about what it means for the row player to have no incentive to deviate. So in this case, when the row player is playing a randomized strategy, what I claim is that the row player has no incentive to deviate if and only if for all i prime the sum from j equals 1 to n of C i prime j yj is at most x transpose Cy.

So what's going on with this expression? There's a reason I'm taking a little bit to build up towards it. You should see that already it Bears some similarity with what's going on here, where the inequality is less than or equal to, but I'm allowing some variation because I'm changing my strategy from my strategy x that has this payoff on the right-hand side to instead being the strategy that just selects i prime.

So this is really calculating what would happen if I play according to x against y versus I decide I'm going to just play strategy. i prime against the strategy yj. So what's going on right here in this expression is it's the same expected value computation. Bob is still playing according to y. And this is the probability that Bob selects strategy j.

But Alice is now deviating to the strategy i prime. So she always selects strategy i prime. And this is her expected payoff. And I want that this expected payoff is less than or equal to what I would get if I just stuck to my pre-announced strategy. So even though you can interpret this as a system of linear inequalities, it's important to be able to translate this back into words to understand what it's modeling. So does this make sense? Yeah? All right.

So let's talk about the column player. And then I'm going to ask you a question, a somewhat tricky question. That'll be good food for later. So similarly, the column player has no incentive to deviate if and only if for all i prime the sum from j equals 1 to n C-- oh, sorry, sum for all j prime, the sum from I equals 1 to m of Cij prime xi is greater than or equal to the payoff.

So this works the same way. Bob is just deviating to a single strategy j prime. The inequality gets flipped because the payments are going to Alice, not to Bob. But this is what it means. At least, This. Is how to turn this definition of a mixed Nash equilibrium into a system of linear inequalities.

Now, let me ask a diagnostic question. So the way that I've defined it, these inequalities, they capture deviating from a randomized strategy, but to a pure strategy. Why not allow more general deviations? So the way to think about this is when I defined the notion of incentive to deviate right here, you started off with a strategy i. And you considered what would happen if you deviate to a strategy i prime. So you started off with a deterministic strategy i. And you deviated to another deterministic strategy i prime.

Here, I define the notion of a mixed Nash equilibrium in terms of the players having randomized strategies. Alice is announcing that she plays according to x. Bob is announcing. He plays according to y. But my notion of deviation was just Alice deviating from her randomized strategy x to a deterministic strategy i prime.

Bob's notion of deviation was deviating from a randomized strategy y to a deterministic strategy j prime. So why didn't I consider deviations to randomized strategies? This is an important concept. And we'll see it a few times in lecture. Anyone have any intuition?

I claim it's the same. So imagine that Alice had a deviation of some other randomized strategy x prime that did better. What I claim is the only way that she has a deviation to an x prime that does better is if there's some fixed deterministic strategy that does better too, because how well she does when I take x and replace it with x prime is just an average of how well she does for each of the different row strategies she can take.

So this is a very important concept. And what I'm claiming here, purely game theoretically, is that if Bob has pre-announced that his strategy is y and he's committed to that, then I don't have an incentive to deviate to-- I have an incentive to deviate to some randomized x prime if and only if there's a fixed deterministic strategy i prime that I have an incentive to deviate from.

So it's not totally obvious. But hopefully, it'll become clearer over the course of the lecture. That's really a statement that has to do with weak duality. So when we talked about linear programming, we talked about weak and strong duality. This is really, as we'll see, the crux of what's going on. And there's a view to think about what weak duality is that happens purely through zero sum games.

So in fact, the starting point of understanding linear programs was actually zero sum games. So it happened with von Neumann's famous work. And it was, of course, famously augmented to not necessarily zero sum games by Nash. But this is really how people first thought about the concepts of weak and strong duality.

So now, I get to tell you about probably my favorite theorem in class. Let me write it up here. This is an amazing result, which we will prove today. But the depth of this result, I think, you appreciate more and more over time because of how many amazing consequences it has.

So this is something called the min max theorem. And it was proven by von Neumann. Let me just write it down as a mathematical statement. There is some value lambda star, which is equal to the max over all x belonging to delta m of the min over all y belonging to delta n of x transpose Cy.

But I claim that this is actually the same thing as switching the max and the min-- so the min over all y in delta n the max over all x in delta m of the same thing, x transpose Cy. So this is my favorite theorem in the entire class we're going to cover. So what this statement is even saying and why it's so deep and powerful, this takes a little bit to digest. So let's go through that right now.

So right now, all I said-- it's called von Neumann's min max theorem-- is that these two things are the same, interchanging the max and the min. And I'm going to call these two expressions 1 and 2. So let's interpret this expression game theoretically, in terms of what the semantics are for these two different expressions.

So what does this result imply? Well, what I claim is that it's really telling us that two different ways to play the game-- actually, three different ways are all equivalent. It doesn't matter. So expression 1, where, I have a max and then a min-- I claim that what's going on with the way that I've ordered the max and the min is that it's a certain semantics for how the game is played, which sounds like it's unfair to the column player because what's going on is that the column player chooses a distribution on strategies, announces it, and the row player best response.

Did I get these backwards? One second. Yeah, I think it's backwards. All right, let's do 2. So what's going on in expression 2 is that first we have this outer min. So first, this-- oh, wait. Let's see. Nope, I had it right the first time. My mistake. All right. It doesn't matter where they're equivalent, but let's just make sure we get the parsing correct.

So imagine that I take that expression. I cover up the max first in my expression 1. So what's going on is that the column player is playing this min strategy. So he's trying to find this strategy y that's a distribution on columns. But now, once he's chosen what that strategy is, once he's chosen what that strategy is, then the max player is allowed to choose the best response.

So even after the column player is fixed and committed to playing y, then the row player chooses the strategy x that maximizes it. I'm going to get this slightly-- yeah, let me just make sure I got these things right. Yeah. All right. I might have gotten it backwards, but that's OK.

And for strategy 2, it'll be the same, but the row player goes first. So the way to think about it is that, if the way that I described the semantics of the game to begin with started with one of the players having to pre-commit to their strategy, a priori that would have sounded like I really changed who had an advantage because now it has to be that Alice is going first or Bob is going first. That person has to commit to what their randomized strategy is.

See, the way to think about it is that, when we just had the notion of a pure Nash equilibrium, then all of a sudden changing the order that the two players play completely changes the meaning of the game. So think about the matching pennies example. If Alice goes first, you definitely want to be Bob because, whatever Alice chooses, you can choose the penny to not match. And then you always win $1.

So when you're committing to deterministic strategies, there's a huge advantage to who commits first to their strategy and who gets to adapt. But what we're saying here is that it actually doesn't matter who goes first, that the best that you can do in either of these two scenarios, where Alice pre-commits to a distribution, then Bob goes or vice versa, is actually the same value. It's the value lambda star.

So in fact, what this min max theorem is really telling us is that these two ways of playing the game, and even the way where we simultaneously choose what our distribution on strategies are, are all completely equivalent. So this may not yet seem like a very powerful result. But I'll give you some more intuition later.

But for right now, let's prove one very powerful thing. So when I talked about a pure Nash equilibrium, that was a very natural solution concept. But unfortunately, we had simple games, like the matching pennies example, that had no pure Nash equilibrium. Now, a mixed Nash equilibrium is more promising because we've augmented the strategy space so that Alice and Bob choose distributions.

What I claim, the corollary of this theorem is that every zero sum game has a mixed Nash equilibrium. So that's the first thing we're going to prove. And then we'll come back and we'll prove the min max theorem based on our understanding of LP duality. All right. So let me state that corollary. And then let's prove it.

I claim every two player zero sum game has a mixed Nash equilibrium. So the main question is, where am I going to get my strategies that are purportedly mixed Nash equilibrium? This part's not so obvious. But it's really going to come from the two expressions right here.

So here's how the proof is going to work. We're going to let lambda star equal the max min, which, of course, is equal to the min max. And now, the main question is, where are my strategy is going to be? So what I'm going to do is I'm going to let x star and y star-- these are my purported mixed Nash equilibrium. They'll be the optimal strategies In 2 and 1, respectively.

So the way to think about it is in these different scenarios. So x star is the answer when the row player is forced to go first. y star is the answer when the column player is forced to go first. And what I'm going to prove is that these pair of strategies really are a mixed Nash equilibrium.

So now, let's put it all together. So we can take x star transpose Cy. And I claim we have the following inequality. So what does this inequality mean in words? So we know that lambda star is the payoff when these two players play x star and y star respectively.

And so what's going on here is that when the row player pre-commits to x star, I claim that, no matter what the column player does for his strategy, that the payoff is at least lambda star. That's what it means for this formulation in 2 to have the value lambda star, is that I can choose my x star strategy. And no matter what happens after I pre-commit to that strategy, I get the value at least lambda star, no matter what that strategy is.

So this is literally just translating the statement of how we're playing the game in scenario 2 into a system of linear inequalities, the same way that we've been doing thus far. I'm going to call this set of inequalities A because I'm going to have to put it together. And similarly, we can do the same thing the other direction. x transpose Cy star is at most lambda star. And this holds for all x that belongs to this distribution on m things.

And this holds because, when the column player has to go first, this is what it means to have a good strategy in scenario 1, is that the column player chooses this distribution y star. And no matter what the row player does, we still get the value at most lambda squared-- lambda star. We're losing at most lambda star dollars. So I'll call this system of linear inequalities B.

And together, what these two things imply is that they imply that x star transpose Cy star is equal to lambda star. Just by putting these two constraints that we have over all y and over all x together, we get exactly what we wanted, that the payoff really is lambda star.

And so what this means, when we put it all together, is that A plus B plus C mean that the expected win loss is lambda star and no worse. So no one can actually do better. So really, just the statement of this min max equals max min is really just the translation of the statement that no one has an incentive to deviate, as long as we're careful about how we get those min max strategies.

And the way we get the min max strategy is we piece it together. One of the strategies comes from this scenario of how we play the game, where it's unfair to the column player. And one comes from this scenario, where it's unfair to the row player instead who has to go first. So are there any questions? That's the first main punch line, is that all zero sum games with two players are guaranteed to have a mixed Nash equilibrium. And that follows just from von Neumann's min max theorem, which we're now going to prove. So any questions? Good? All right. So now, let's prove the min max theorem.

The first thing I'm going to do, before I prove the min max theorem, is I'm going to tell you a lemma that makes precise something that I asserted earlier, that the notion of deviation-- you can, without loss of generality, just deviate to pure strategies. Let's make this precise. And then we'll prove it. And we'll see how it's really weak duality, types of arguments and action.

So here's my key helper lemma we're going to use to prove the min max theorem. I claim that the constraint-- or rather, the system of constraints x transpose C xy is at least lambda. And this is true for all y in delta n. Let me call this the constraint system 3. I claim it is equivalent to the constraint x transpose C Ej is at least lambda. And this is true for all j belonging to n. And I'm going to call this constraint system 4.

So this just makes precise the thing which I asserted earlier. I'm just saying that, when we have a system of constraints, this is an extraordinarily complicated system of constraints because there are infinitely many constraints, one for every choice of y that's a distribution on n things.

And I want the constraint that, no matter what you choose for y, my payoff when I play x is at least lambda. That's what this constraint is saying. No matter what randomized strategy you try, I get at least lambda. I claim that this system of constraints that's captured by this equation 3 is equivalent to this much simpler system of constraints that, when I play x, I just have to consider strategies for you that are deterministic strategies.

So here, Ej is the vector that has 1, 1 in the j-th coordinate and the rest of it is 0's. So what this corresponds to is instead of y being a general distribution, it's now a fixed choice of playing strategy j. And I just need to ensure that against every one of your individual strategies I get at least lambda.

So this is the crux of what I was saying earlier, that I only need to consider deviations which are deterministic strategies. It's really the statement that the system of linear inequalities are actually equivalent to each other. You can derive one from the other. So let's prove this fact. The proof is very simple.

So one direction is obvious because, whenever I satisfy the constraints in 3, I better satisfy the constraints in 4 because this is just a subset of them. It's a very particular set of distributions I look at. But the more interesting direction is when I satisfy all the constraints in 4, how exactly can I prove that I satisfy the constraints in 3?

So here, the key point is that if I satisfy all the constraints in 4 and I want to prove I satisfy the constraints in 3, what I'm going to do is I can take an arbitrary constraint y that belongs to the set 3. And what I can do is I can take this vector. And I can break it up into the sum of simpler vectors. So it's the same thing as the sum from j equals 1 to n of yj Ej.

Remember, Ej is just this placeholder that has a 1 in the j-th coordinate and has 0's everywhere else. So this is just a simple equality, breaking up this vector y into each one of the pieces for each one of its coordinates. And now, I'm in business because, when I consider an arbitrary constraint of the form x transpose Cy is at least lambda, I claim that I can derive it as follows.

I can take the sum from J equals 1 to n of yj x transpose C Ej. Let me just put the inequality inside. So the way to think about it is that if I take this vector y and apply this substitution and use linearity, I can just pull out all the yj's and the summation. And now, what's going on is that this inequality right here is just a weighted linear combination, non-negative linear combination, of the inequalities that arise in 3.

So each one of these inequalities shows up in 3 because it corresponds to an Ej. And this tells me that I can actually go from satisfying these equations in 3, add them up in order to derive any constraint I want that arises in the constraint system 3. So this should remind you of weak duality because what's going on when you write the dual to an LP is you're really asking, when I have a maximization problem, how can I set up a companion LP that's a minimization problem that derives an upper bound on the largest value of the max I can get?

The way that you form the dual is really by taking the inequalities in your original system and adding them up with the appropriate variables. So your dual is just asking what's the choice of how to add up those inequalities to derive the best constraint. That's exactly how you prove weak duality. And it's really the crux of what's going on from a game theoretic standpoint here, too.

So are there any questions? Make sense? Once we have this lemma, now let's prove the min max theorem. And because we already have all of the tools of linear programming duality, we're just going to prove it by appealing to LP duality. So let's prove the min max theorem using this helper lemma.

So what we're going to do is we're going to write down an LP for playing in scenario 1. So the proof strategy is very simple for this. What we're going to do is I told you that the min max theorem is really about saying that two different scenarios for playing the game are equivalent. We're going to write down an LP that captures each of these scenarios, that capture scenario 1 and that captures scenario 2.

And what we're going to observe is that those two LPs are actually going to be dual to each other. So that's the way that the proof of the min max theorem goes. So I'll suggestively write p for primal. So we want to maximize lambda such that-- one second-- ah, this would actually be scenario 2. So let's do scenario 2. Sorry for getting it backwards.

All right. So for scenario 2, we want to maximize lambda subject to the constraint that x transpose C Ej is at least lambda. This should be true for all j in my strategy set n. I want that the sum of the xi's from i equals 1 to m is equal to 1 and all of the xi's are non-negative for all i.

So really, what's going on is I'm trying to maximize this value lambda that's a lower bound on how well I do. But part of the action is really the choice of what the x is. So I want to find a strategy x, which is a distribution on the strategies, with the property that no matter what the other player plays in response-- I've chosen x. No matter what the other player plays in response, for every strategy j, I will guarantee you that, in average, I win at least lambda.

And now, I don't just want to find an arbitrary x that satisfies this condition. But I want to make lambda as large as possible. So I want to maximize my safe value so that there exists a strategy x so that, no matter what you do, I get at least this value lambda in expectation. And this is the same thing as scenario 2 because, once the row player has committed to playing this strategy x, then what would the column player do? They would choose the j that minimizes this quantity, so they lose the least amount of money, which is the same thing as asking that it's for always at least lambda.

So that's my primal. And now, I can do the same thing in scenario 1. So I'll suggestively write d for dual here. And here, what I would do is I would minimize over all lambda. I would want the constraint that Ei transpose Cy is at most lambda for all i in m. And of course, I want the constraints that my vector y is a distribution-- so the sum from j equals 1 to n of yj equals 1. All the yj's are non-negative on each of the coordinates.

And so here, it's the same type of thing, but in reverse. My column player is going first. My column player is choosing a strategy so that no matter what the row player does, you can guarantee that an expectation I lose at most lambda dollars. And I want my lambda, my upper bound, on how much I lose in expectation, to be as small as possible. It could even be negative if I win dollars in expectation.

So these are the two LPs. And the proof of the min max theorem now is very simple because what you can do is you can check that P and D are duals to each other. And moreover, both are feasible. Why? Because I can just choose any x that's a distribution, any y that's a distribution. That'll be some finite value for lambda in both of these two scenarios. So feasibility is not a problem. And once you have that both are feasible, you know that strong duality holds.

So we know that their optimal values are equal. So this gives us von Neumann's min max theorem. p is equal to max min, which is equal to d, which is the same thing as min max. So this is also not just a cool application of LP duality, but it really gives you a better way to think about LP duality.

| what is weak duality? Well, one way to think about what that concept means is, when I'm in scenario p and I'm playing according to this scenario 2, one way I can prove for you that you can't get a very large value lambda is if I just have to exhibit a good strategy for the other player.

Once I have a feasible solution to this linear program, that tells me how the column player has to play. And that tells me that my row player can't do all that well. So these two things act as certificates of each other. Whenever I have a feasible solution to one side, then it implies something about the other side not being too good. And that's really the crux of what weak duality is all about. And strong duality just tells us that, actually, the best certificate we can get for p from d is actually equal to the optimal value that we can obtain. All right. So any questions? Make sense? All right.

So we covered a lot of ground already. So we talked about zero sum games. We talked about pure Nash equilibrium and how sometimes they don't exist. We extended our solution concept to mixed Nash equilibrium. And once we had mixed Nash equilibrium, then there were a lot of basic questions like, do they always exist?

We proved that they do exist by courtesy of the von Neumann min max theorem that told us where x star and y star came from. And then we proved the min max theorem just as a consequence of LP duality. So one of the amazing things-- I told you that this is one of my favorite theorems.

So actually, first let me tell you the-- let's go back to the example we had earlier, where I had you guys play that game. So we can ask the same questions that I asked back then. So when I wrote down that C matrix-- let's just write it down one more time. So I had this payoff matrix C that was 3 by 3-- 2, minus 1, 3, 2, 2, minus 3, minus 1, 2, minus 3.

So we can ask the same types of questions. Who's in the better position? Maybe you still wouldn't know the answer. But how could I find out who's in the better position? So how would we answer this question now, now that we have all this technology? What should I try and compute? Yeah?

**STUDENT:** I guess the optimal solution [INAUDIBLE].

**ANKUR MOITRA:** Yeah, that's right. So one thing I can do is I can compute the game value, lambda star. Lambda star, which is called the game value. Remember, the game value captures, if both players played optimally, who would win how much in expectation?

So one way I can answer this question of who's in the better position is I can compute the game value, lambda star. Now, in simple examples like the matching pennies, computing the game value is easy because it's 0. We know what the best strategy is. Every player plays randomly. And then everyone wins nothing in expectation.

But now, for more complicated games, where it's not so obvious who you would want to be, I can just feed this to my linear programming solver. And I can compute whatever the lambda star is. So in fact, the thing you can show is that lambda star is actually 1/3, 1/3 of $1.

So what this means is that-- so who should I be, the row or the column player? The row player. So if we'd played with real money, then you actually would have won money if you played according to the min max optimal strategy. Maybe sometimes you get unlucky and you lose some money. But it's not actually this game of whack-a-mole where one player plays a good strategy and the other player switches to try and get a better strategy and do better in response.

In the beginning, you lost money. And then you quickly came back. But we know what this game of whack-a-mole converges to, at least if both players play rationally, because both players should really just play their mixed Nash equilibrium. And you can check that the x star that does this-- let me just write it down.

So you can actually solve this just by hand or plugging it into an LP solver. And what you'll end up getting is that x star is equal to 5/9, 0, 4/9. And y star is equal to 0, 2/3, and 1/3. And we can check that x star transpose C is equal to this vector 2/3, 1/3, 1/3. So that's my proof that the game value is at least 1/3, is when I play that randomized strategy when I am the row player, well, this vector right here just describes whatever the column player responds to, what would be my expected value.

The column player should never play this strategy. And they'll only play among these two strategies. And I've guaranteed myself that I'll win at least 1/3 of $1. So did our two volunteers play optimally? So you can think about it. What this tells us is that, in the optimal strategy, you should actually never play, if you're the row player, strategy 2. But that happened.

If you were the column player, you should never play strategy 1. I don't mean to pick on you. It's a hard game. That's why I picked it, was because it's not obvious. And this is really what's powerful about this solution concept, is we took something that we had no idea how to play and we just figured out how to play it. We figured out how to play it optimally.

Now, in fact, let me ask you an even tougher question because really, my goal is not just to teach you about zero sum games, but to give it to you as a vantage point for how to think about the key concepts from linear programming in a more intuitive way, at least in a way I like to.

So one thing you'll notice, I take x star and I take x star transpose times C. And I get this vector 2/3, 1/3, 1/3. Now, only these things are tight with the game value. These are the only things that match what I claim my lambda star is. What concept from linear programming theory tells me why I must have a 0 here? This is a very tricky question.

We already covered. Weak duality. Weak duality in the context of game theory was how these two different scenarios for playing the game-- one which seems to advantage the row player, one which seems to advantage the column player-- figuring out how to play one of those scenarios bounds what's possible in the other scenario. That was weak duality.

I claim that when I take x star transpose times C, and x star and y star are my purported optimal solutions, and I look at where I have tight inequalities that these values for the payoffs are actually equal to lambda star, the game value, that anything that's not tight has to have a matching zero. What concept? I hope Peter covered this. Yeah?

**STUDENT:** Is it basic?

**ANKUR MOITRA:** It's related, very close. What about complementary slackness? So this is complementary slackness because whenever I have an inequality that's not tight, like this one right here-- this is not a tight inequality. It's not equal to the lambda star-- I better have the corresponding variable is equal to 0.

The only way that I can have strong duality hold is as long as that's true. And if instead I took C times y star and I wrote it out, I would get 1/3, 1/3, 1/3. So I don't get any constraints. But if instead I had some value was not equal to 1/3, then the corresponding thing in x had better be equal to 0, too.

See, the way to think about complementary slackness is that if this is my purported optimal solution, if x and y star are my purported optimal pair, then as soon as y star has some positive mass on this outcome, then I actually am winning more in expectation than 1/3. So my game value cannot be the correct answer. So really, also you can think about complementary slackness as this constraint about the compatibility of their two strategies and the associated vectors of expected payoff.

So let me briefly mention-- this is way beyond the scope of this class. But one of the reasons this is my favorite result, at least in the things I teach in undergraduate classes, is actually my PhD thesis started in some way from this. So I had a result that I wanted to show in graph theory, which I'm not going to explain what it is. But it's a very complicated and surprising result.

So I remember I went into my advisor's office in Akamai. And then I told him I proved this amazing theorem about all kinds of graphs, about a structural theorem about them. And he said, there's no way you proved that. That doesn't even sound true to me. So he started writing some concrete examples. What would your sparsifier look like for this graph? I couldn't do it.

He gave me another example. I couldn't do it. He gave me another example. I think he started to feel bad for me because he seemed to get the hint that maybe I hadn't actually proven it. And he said, how could you prove this if you can't construct any of them on the graphs I give you?

And the way that I proved it was, I actually appealed to a gigantic zero sum game, where I was playing one of the players, the row player, and he was playing the column player. So I didn't have to prove that such things existed by actually constructing them myself. I just had to show that any strategy he had for disproving their existence didn't work.

And that gives you a very indirect way to prove existence of pretty exotic objects, things which are hard to wrap your head around otherwise. So even though you might not see this until much later if you see it, these types of things are really under the hood of a huge amount of combinatorial optimization. It's not just about zero sum games I can write down because sometimes it's about truly gigantic zero sum games that are just exponentially large or impossibly large to wrap your head around.

And instead of this whole idea about playing the game in two different scenarios, instead of me having to prove to him that they existed, I just had to prove that he couldn't prove me wrong. This gives you a great way to prove all kinds of exotic things that otherwise sound counterintuitive and like they should be false.

So what's really amazing is when you take a lot of results across theoretical computer science that sound false to you and you don't have the intuition-- it sounds like they shouldn't exist-- it turns out that, when you pull on the thread, a lot of times this is what's under the hood, many times.