[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: So then let's get started with our second-to-last topic today, which is going to be something called cheap talk. I think it's good to have some context for this.

So far, we've looked at signaling games. And so let's recall that when we looked at signaling games, people with private information were conveying information to some other party through costly actions.

So under signaling, info is conveyed through costly or we might also say payoff-relevant actions. So one way to convey information about how smart you are or how good of an employee you are is to go to college.

And the whole point of that model was that going to college is costly and differentially costly for different kinds of people. Or when a firm was conveying information about the quality of an item, they were setting a price. And maybe a price isn't directly costly, but it is payoff relevant, because the price I set affects how much money I get if you actually accept my offer.

But in the real world, we see a lot of information conveyed not through costly actions, but just through speech, through communication. We send messages to people. We use words. We talk to people.

And the question is, when there's a conflict of interest, can we also convey information through costless actions? So today, the difference is we're going to look at conveying information through costless-- and sometimes, instead of calling these actions, you might call them messages. Because we usually think of speech or words or messages as being costly.

Now, there are certainly some contexts where this isn't going to work. If we think of the job market signaling example, the firm could ask you how good of an employee are you. Well, everyone's going to say they're a good employee.

So there are some contexts where we really need a costly action in order to make it credible. The information that we're going to convey. There are other situations where it's clear that communication is credible.

If our interests are completely aligned. Then of course, we can communicate freely and that's not going to be difficult. So the point of the model today is going to show that sometimes, there will be a misalignment of interests.

There will be a conflict of interests. Yet nevertheless, some kind of communication can be credible, in contrast to the model of telling the employer that you're a good employee. So that's going to be the model we're going to look at today.

So this model is based on a paper by Crawford and Sobel in 1982. So we're getting closer to modern topics. And we're going to think of-- this is a game between a sender. We call them sender and receiver.

But these are abstract labels. I want you to think of the sender as an advisor and the receiver as a decision maker. And the presidential advisor might collect information and make a recommendation to the president about what tax policy should be, or to the Fed about what the interest rate should be, or, how much we should spend on the military.

And the idea is that the sender, the advisor, has more information than the decision maker, because the decision maker is very busy. The CEO can't know about everything. They talk to a consultant or a researcher who's going to be very informed about this particular issue.

So how do we model this? Well, the sender is going to privately observe what We're going to call maybe a state. Or this could also be called their type t in 0, 1.

So we're going to think of this in our economic example. How good is the economy? What's the ideal amount of spending we should spend on the military? What's the ideal tax rate?

So this makes sense for any decision where there's this one-dimensional choice that we have to make. We can think of this usually as the state of the world. But we'll also call it the sender's type, because this is private information that the sender has.

And then the sender, after observing their type t, is going to send a message m-- we'll just call it, in some abstract set, big M-- to the receiver. And then the receiver is going to observe M but not t.

The whole point is if the receiver, the decision maker, understood the state of the world, knew what the right decision was, they wouldn't need help from the advisor anyway.

So they don't directly observe t, but they do observe this message that's sent by the sender. And then after that, they choose an action or decision which we'll call y. And that's going to be a real number.

So maybe I'll say-- I don't know if I've used this notation before. This is just the set of real numbers. Now, whenever we talk about something being privately observed, we need to specify the distribution from which this comes. And we're going to assume that the state, just to make things easy, is uniform.

So just to be clear, the sender knows what the state is. They know their type. The uniform distribution is relevant to the receiver. The receiver believes that this state is uniformly distributed, but they don't know what the actual value of the state is. I'll say uniformly distributed.

And now let me specify payoffs. Of course, this is a very stylized model, but I think it captures a lot of policy disagreements that we have. So the sender, their payoff u s depends on the decision y that the receiver takes and also t, the state of the world.

And same for the receiver. And what we're going to imagine for the receiver, to make things simple, is the receiver just wants to match their decision with the state of the world.

So we could think of the state of the world might be some measure of how close the economy is to a recession. And y might be how much we're planning to cut interest rates.

Or t might be the severity or the risk of some military conflict, and y might be the spending level that we spend on that conflict. And the idea is, these things move together. The higher the state, the higher the ideal decision for the receiver.

So we'll formally model this with the quadratic loss function that we said before. So this means what the receiver would love to do is have the decision exactly match the state. So we can think of the state as basically measuring what would be the ideal decision for the receiver if they could observe the state.

And then the receiver experiences a loss that's equal to the difference between their decision and the state squared. So how far off the decision is from the state. Yes.

AUDIENCE: So basically, the decision is just a reflection of what the receiver believes that the state is?

PROFESSOR: It's true that their preferences over the decision depend on their beliefs about the state.

AUDIENCE: But [INAUDIBLE] exactly match the state, right? So it should be-- so is the decision supposed to be a state?

PROFESSOR: No, I wouldn't. I think maybe one way of interpreting this is, OK, if we want to think of a richer model, probably the state is really a very rich multidimensional object.

But we're going to denote the state-- the aspect of the state that matters is what the ideal decision is in that state. So I think maybe the better interpretation is, when we write t, what we really mean is, what is the receiver's ideal decision given the state of the world?

So maybe the economy has many dimensions. But think of t saying the ideal interest rate for the Fed to set is 3% and y being the interest rate the Fed actually sets. Does that make sense?

AUDIENCE: Yeah.

PROFESSOR: OK, great. So now what about the sender? Well, if the sender and receiver have the same preferences, this would be a really easy problem. Of course the advisor is going to just reveal the state and allow the receiver to make a good decision.

But often, advisors and decision makers have misaligned interests. And the way we're going to capture that here is we're going to say the sender's preferences is like this, where beta-- I'll put it down here-- is greater than 0.

So here, beta is called the bias. This is the bias that the sender or the advisor has relative to the decision maker. So notice that, well, the receiver always wants to match the decision with the state.

The sender would like to match the decision not with the state, but with the state plus this bias beta. So notice some characteristics of this misalignment of interest.

Both the sender and the receiver agree that when the state is higher, a higher decision is warranted. What they disagree about is the exact value of the decision. So whatever the state is, the sender always thinks the decision should be a little bit higher than the receiver thinks.

And I think this is consistent with a lot of policy disagreements. Let's take the example of the Fed. There are some people who are more hawkish or more dovish on inflation. So some people always tend to want lower interest rates than someone else.

But everyone agrees that when the economy is in worse shape, we should tend to cut interest rates, and when the economy is doing well, maybe we should raise interest rates. So they all agree about the directional effect of the state on the ideal decision. They just disagree slightly about what the level is.

And the sender always wants a decision that's higher than what the receiver wants exactly by this amount beta. So the larger beta is, the greater the disagreement between the sender and the receiver. OK, so here's our model. Any questions about the model so far? Yes.

AUDIENCE: Does beta just depend on who the sender is? It's not like distributed, it's just like [INAUDIBLE]?

PROFESSOR: Good question. So you could imagine a richer model where the decision maker doesn't know the political biases of the advisor. And that's actually an interesting extension that people have looked at.

But for today, we're going to assume beta is commonly known. So I hire this advisor. I know that my advisor always tends to have this policy preference. And this is known. What I don't know is the true state of the world. Good question. Any other questions?

OK, so it's a game. Let's first think about what strategies are in this game. So the sender, they observe the state of the world or they observe their type t.

And they choose what message to send to the receiver. So a strategy for the sender is a function m from 0, 1 to this abstract set of messages M. Maybe we'll go into here just to understand what goes on.

So what this says is if I see that the state is, say, 0.3, this is the message I send, m of 0.3. If I see that the state is 0.7, then I might send a different message, m of 0.7.

Now, the receiver's strategy, well, what do they do? They see the message. And based on the message that they see, they choose a decision to make. So for the receiver, a strategy is a function y from M to we'll say R. Just to understand what's going on, I'll write it like this.

So the receiver says, if I get a message m from my advisor, this is the decision I'm going to take, y of m. And the solution concept we're going to work with is going to be PBE, as before-- Perfect Bayesian Equilibrium.

Now, let's start by analyzing how things work for the receiver. So with perfect Bayesian equilibrium, whatever message the receiver sees, they're going to update their belief and then behave optimally given their updated belief.

So let's just remember that whenever we work with perfect Bayesian equilibrium, we have to distinguish between the messages that occur on path and the messages that don't occur on path.

So a message m is on path if, in a given strategy profile, there's some state of the world where the sender actually sends that message. And a message is off path if that message is not supposed to be sent in any state of the world.

So if the receiver sees an on-path message m, they apply Bayes' rule. And then using Bayes' rule, they have an updated belief about the state of the world t.

But because of the nature of our preference relation, the nature of our utility function, the receiver always wants to take the decision that's equal to their updated expectation of the state. So they apply Bayes' rule. They can now compute maybe I'll say E of t given m.

What does this mean? It says, I observe m. I update my beliefs about the state of the world. And I'm now going to compute what I think the expected value of the state of the world is given my updated beliefs.

And it's always going to be optimal for them to take this as their decision. Compute that, and in fact, that's going to be their decision. They're going to update their beliefs about the state, compute an expectation and set their decision equal to that.

Off pat, well, the beliefs over 0, 1 are unconstrained because we can't apply Bayes' rule. So off path, the receiver can believe anything about the state between 0 and 1.

So the key question is, what decisions could the receiver take if they see an off-path message m? What are the range of decisions that are consistent with some belief? Yeah.

AUDIENCE: Just the real numbers.

PROFESSOR: So we might think it's all the real numbers. But is there any belief that's consistent with taking decision y equals 2? We know t lives in 0, 1. So the point is-- and this is kind of a subtle point-- the point is, I can form any belief I want about the state. But we know the state lies between 0 and 1.

So the most pessimistic belief I could have is the state is definitely 0. And the most optimistic-- or here it's not clear what's optimistic, but the highest belief I could have is the state is certainly 1.

And in those two cases I would take the decision y equals 0 and y equals 1. So I'm always going to take some decision. y that's always going to be between 0 and 1, because those are the decisions that are consistent with some belief about the state.

Any decision in here can be justified by some belief about the state. I could take decision 1/2 because I think the state is certain to be 1/2. Or maybe I think it's uniform. Those are all going to be OK.

It's going to turn out, as we'll see, that the off-path beliefs are not going to play as big of a role in this model as they did in the signaling model. But we'll see that a bit later.

So let's start with, just to build intuition, a really simple equilibrium. And the simplest equilibrium would be a no-communication equilibrium. Now, when I say no communication, what I mean is no information is actually conveyed.

Of course, they're actually going to send a message because they have to. But one thing the advisor could do is they could be a broken record. They could just always send the same message. And if they always send the same message, no matter the state of the world, then that message doesn't convey any information about the state.

So let's look at this kind of boring equilibrium. Well, what is-- let's fix some message I'll call m* in m. This is going to be the message that the sender always sends, no matter the state of the world.

So the sender's strategy is m of t equals m* for all t in 01. So going to our economic example, if they see the economy's great, they say cut interest rates. If they see the economy is bad, they say cut interest rates. And therefore, what they say conveys no information about the state.

Now let's try to compute what the receiver is going to do in this case. So let's now look at the receiver. And what happens if the receiver sees the message m*? What beliefs does the receiver form about the state if they get this message m*? Yeah.

AUDIENCE: Were they just going to be left with the prior, which is that t is distributed over 0, 1?

PROFESSOR: Exactly. They'd just be left with the prior, because they know whatever the state is, they're going to get this message. This hasn't told them anything. If you always say the same thing, you haven't revealed anything to me.

So my beliefs are the same as the prior. So if I see m*, maybe we could say retains prior beliefs. That is, I still think the state is uniformly distributed over 0, 1.

And the average state is therefore 1/2. So the best decision for me is to take y equals 1/2. So it retains prior beliefs and they take-- so y of m* equals 1/2.

So this looks good for equilibrium. But we still have to worry. Well, we have to say, what does the receiver do if they see a message m That's not equal to m*?

They don't really know what to do in this case, because they were expecting to always see the message m*. They can't really apply Bayes' rule here because they'd be dividing by 0. So there's a lot of flexibility.

You might say, oh, well, let's just choose y arbitrarily. But what could go wrong if we just chose y arbitrarily here for the receiver? Let's say we chose y equals 1/3 or something. What could go wrong here? Yeah.

AUDIENCE: Maybe the sender will-- like, that would be a profitable deviation.

PROFESSOR: Exactly right. So whenever we're choosing the off-path behavior by the receiver, we have to keep in mind that we don't want to choose behavior that will introduce a profitable deviation for the sender.

In this equilibrium, the sender is only supposed to send the message m*. So we want to make sure that the sender never wants to deviate to send any message other than m*. So how can we be sure? How could we choose y of m to be sure that the sender never wants to deviate? Yeah.

AUDIENCE: You could make it so that the difference between-- you get the same utility from getting m* and r.

PROFESSOR: Exactly right. And how could we make sure they get the same utility?

AUDIENCE: I don't know exactly, but it's depending on beta, right?

PROFESSOR: I think you're on the right track. There's even a simpler answer. So you're exactly right. We want to make sure that if they deviate to this, the sender gets exactly the same utility. But what could it be? Yeah. It could just be 1/2.

So what they could do is they could say, look, if I see any other message, I'm going to do exactly the same thing. I'm just going to choose ym equals 1/2. I'm going to maintain my prior belief.

So let's understand the structure of this no-communication equilibrium. The receiver says, whatever message you send me I'm going to do the same thing. I'm going to completely ignore you. And the sender says, well, if you're going to ignore what I say, it doesn't matter what I say. So I might as well always send the message m*.

And this kind of sometimes goes over the name-- this is sometimes called babbling. So in this equilibrium, the receiver is treating his advisor's words as just babbles. I don't care what you say.

It's not going to change my belief. I'm just going to ignore what you say and do the same thing. And therefore, the advisor says, well, I might as well always send this message m*.

So indeed, this satisfies-- let's just check what we would have to check to make sure this equilibrium. We'd have to check sender optimality, receiver optimality, and belief consistency.

I've been a little imprecise. I haven't really written what the beliefs are here. But naturally, the beliefs could be, I just think the state is uniformly distributed if I see any other message. And then we can see that each of these three conditions is going to be satisfied.

This is not a very interesting equilibrium, but it's just good to see there's always an equilibrium like this, where our words just don't have any meaning And no one listens.

Well, that's one extreme. Let's see if we could do better. We might want to have perfect communication. So what about an equilibrium where there's perfect communication?

So you might think, well, if the advisor and the decision maker have a pretty good relationship, what should happen is the advisor should just say this is what the state is. So could there be an equilibrium where the message m of t just equals t?

So the advisor says, this is the state, make a decision. Any thoughts? Could we make this an equilibrium? Yeah, maybe?

AUDIENCE: It would not be an equilibrium because the receiver would just act with y equals t as a result. And then the sender is losing some amount because y does not equal t plus beta.

PROFESSOR: Right. Let's work through it. If we tried to-- this is not going to work, but let's go ahead. If we tried to have an equilibrium where this was the messaging strategy of the sender, well, this looks great for the receiver.

What is the receiver y of t going to be? Let's be clear here. The receiver's strategy is not a function of the state. It's a function of the message. But here, the message is the same as the state. So that's why I'm writing t.

Or actually, maybe I'll still write m. Eh, torn about which notation to use. It doesn't make a difference. I'll write m. In this case, when they see a message m, they're going to say, well, that means the state t equals m.

So their best decision is just going to be m, right? But now let's put yourself in the shoes of the sender. And let's say you see that the state is t. If you were the sender, what message would you send to the decision maker?

AUDIENCE: t plus beta.

PROFESSOR: Exactly. So you know the state is t. Well, why don't you just say, look, you're going to tell your advisor m equals t plus beta? And I think we see this all the time.

If someone's more of an interventionist and they know that the president is less interventionist and they want them to intervene in the conflict, they don't say, I'm interventionist. They say, oh, this conflict is just really serious. It's really important you intervene.

I'm not saying it because I'm interventionist. It's just the nature of the conflict means we really have to intervene. Or someone who always wants to cut interest rates isn't going to say, oh, let's cut interest rates because I love cutting interest rates.

They're going to say, inflation's not really a problem. We're not really having any inflation. I think this is what you should do.

So the sender has an incentive to misrepresent what the state is in order to get their ideal decision. Because they know if the receiver believes the state is actually t plus beta, then the receiver will take decision t plus beta, and that will give the sender exactly what they want.

So we see that indeed, this cannot be an equilibrium, because the sender has a profitable deviation, and then the receiver won't trust the messages, and everything breaks down. So not an equilibrium.

But now we have a bit of a puzzle. We definitely see communication in practice between people who have misaligned interests. We don't just see this no-communication equilibrium. So the question is, well, if we can't get a perfect communication equilibrium, maybe it's possible to have an equilibrium where we have some partial communication.

We don't perfectly reveal the state, but we provide some valuable information about the state. And that's going to be the key idea of this paper. So let's our goal is to find, well, maybe not perfectly informative equilibria, but partially informative equilibrium.

Let's start with a really simple case. I guess the simplest possible message you could send about some real valued state is you could say the state's either high or low. You could just say it's either above a threshold or below a threshold. So let's see if we can work with that.

So the state is here between 0 and 1. And let's say we choose some threshold. We'll call this threshold t1. And what's the sender going to do? If they see the state is below t1, they're going to send a message.

Maybe we'll call the message L. They're going to say the state is low. The state is below T1. And if they see the state is above T1, they're going to send a message H saying the state is high.

And we want to see if this could be an equilibrium. You might worry because you might think, well, wait a second. Isn't the sender going to want to lie and pretend the state is high when sometimes it's not in order to get their higher preferred decision?

That's a good intuition, but it turns out we can get this to work, which I think is maybe a bit surprising. So let's go through this. So what is y of L going to be and y of H going to be?

So t1 Is a parameter that we're going to solve for. So we're going to look for an equilibrium that has this form for some parameter t1. And this is what we're going to figure out.

So what happens if the receiver sees the message L? Well, the receiver knows that the state is somewhere between 0 and T1. And they know it's uniformly distributed between 0 and T1.

So what do they think is the average expected? What do they think is the expected state when they see message L? Yeah.

AUDIENCE: 1/2 of t1.

PROFESSOR: 1/2 of t1. Because they're going to say, look, all you've told me is the state's here. It's uniformly distributed. So on average, it's the midpoint of this interval, which is t1 over 2.

Conversely, if the receiver gets the message high, they're going to say, OK, the state is uniformly distributed between t1 and 1. So my updated expectation of the state is just the midpoint of this interval, which is the average of t1 and 1. So this is just going to be T1 plus 1 over 2.

Now, we're not done yet. We still have to worry what happens if they get a message that's neither L nor H, right?

But there's kind of a trick that we can always use that goes back to what we did over here. The trick was, if we see a message that we're not supposed to see, we can just treat it as if it were a message we were supposed to see. And that's never going to be a profitable deviation.

So in this case, the receiver can say, look, if you send me any message other than L or H, I'm just going to treat that as if you had sent me L. So it's going to be t1 over 2.

So as if m equals L. And then there's never a reason for the sender to deviate to any message other than L or H, because if they send a message other than L or H, it's just as if they're sending the message L. So they might as well send the message L to begin with.

And this is a common kind of trick that's used in cheap talk games. I want to emphasize, I don't want to create confusion with signaling. So I want to say this is a trick that doesn't work with signaling.

Let's just understand this, because I don't want to create confusion going back to the signaling game. What if, in the employment context, the employer said, OK, you're supposed to get this level of education? If you don't go to school, I'll just treat it as if you went to school. And that shouldn't be a profitable deviation.

That's not going to work, because they'd much rather not go to school and form the same beliefs. So crucial reason this trick works is if I can send a message L and message m and get the same decision, I'm indifferent between those two things.

But if the message itself were costly, then I wouldn't necessarily be indifferent, because I'd have to take into account not only the decision it induced, but the cost of the message I'm sending.

So this is a crucial difference between a model with payoff-irrelevant messages and a model where the actions I'm choosing are payoff-relevant. I that's a bit of a subtle point. So any questions on that difference? Yes.

AUDIENCE: Could you repeat the example?

PROFESSOR: The employee example? Yeah. So in the employee example, let's say we're doing the job market signaling. And everyone is supposed to get-- let's look at the pooling equilibrium, where everyone goes to college and gets a certain wage.

And the employer says, OK, if you choose any other level of education, I'm going to treat it as if you'd gone to college. And I'll pay you the same wage as if you'd gone to college.

Well now, the student is going to say, wait a second. If I can instead of going to college not go to college and get the same wage, I'm strictly better off. And the difference here is, even though those two actions induce the same decision by the receiver, the actions themselves have different costs. And therefore, the student would rather not go to college.

Here, the message L and the message m induce the same decision. But in this case, the sender is actually indifferent between those two things because the message itself doesn't directly have a cost. Does that clear things up? Yeah, great. OK.

And actually, that's why what I said earlier-- off-path beliefs and belief consistency don't really make a big difference in cheap talk games. And everything we do is basically the same as just doing Nash equilibrium.

But that's a subtle point. That's not so important. So here we are. Let's see if this is going to work out. We've specified the receiver strategy. And now we have to check that the sender is willing to send these messages as they're supposed to.

So what's the sender supposed to do? If they see any state here, they send message L. And if they see any state here, they send message H.

So how can that be optimal for them? Well, what must be true if the state is exactly t1? So let's ask this. What must be true about the sender if the state is exactly t1? Yes.

AUDIENCE: The sender is indifferent between sending L and sending H.

PROFESSOR: They have to be indifferent. Because if the state were just a little bit to the left, they'd want to send L. And if it's a little bit to the right, they'd want to send m.

So the only way they can have those preferences is if at t1, they're exactly indifferent between sending the message L and the message H. So the idea is here, I'm going to be exactly indifferent. And if the state is anywhere to the left, I want to send L and anywhere to the right I want to send H.

So what we need is we must have indifference. So let's work this out. Well, that means when the state is t1-- maybe we'll write it graphically. What we must have is y of L.

Maybe we'll put negatives here. It doesn't matter. So if I know the state is t1, if I send the message L, the receiver is going to take decision y of L. And this is going to be my utility.

If I send the message H, the receiver is going to take the decision yH and this is going to be my utility. So in order for me to be indifferent between these two things, I need these two formulas to agree.

But let's look at it graphically because I think that's a bit clearer. Well, graphically, we have a quadratic loss function. I'm the sender.

I know the state is t1. What I'm going to graph is, well, what does my utility look like as a function of the decision y that's taken? Well, my utility is going to look something like this.

Here it's actually going to be 0. So if the decision y is exactly t1 plus beta, well, that's the best thing that I can get. The decision exactly matches t1 plus beta.

The difference between the decision and t1 plus beta is 0. My loss is 0. And that's the best possible thing I could get. If the decision is higher than t1 plus beta, well, I'm going to experience a loss that's quadratic. And if it's lower than t1 plus beta, I'm also going to experience a loss that's quadratic.

Now, the nice thing about the quadratic is it's symmetric. So what we need-- well, we need the picture to look something like this. We need the sender to be indifferent between the decision yL and the decision yH.

And what that means is the utility function has to take the same value at these two decisions. It crosses through the same horizontal line. But because of symmetry, this is only going to happen if t1 plus beta is exactly halfway between these two points.

If t1 plus beta was closer to yL, then the sender would rather induce the decision yL. And if t1 plus beta was closer to yH, they'd rather induce the decision yH.

So we could do it algebraically. But I think it's clearer to do it this way. So what we see is t1 plus beta must be halfway between these two things. So it must be yL plus y of H over 2.

Now let's do just a little bit of algebra. Maybe let's multiply both sides by 2 to make it a little easier. We have 2t1 plus 2 beta is equal to y of L plus y of H. And let's just plug these in.

So y of L is t1 over 2. y of H is t1 plus 1 over 2. And that's t1 plus 1/2. I just put the t1 over 2 and t1 over two together. And now let's do a final simplification. Get t1 equals 1/2 minus 2 beta.

So let's see if we can visualize what this looks like. If beta is 0, let's start with that case. That's the easiest case.

If beta is 0, then t1 equals 1/2. So that means that the sender just says, is the state above average or below average? But if beta gets larger, then t1 has to be below 1/2.

So this interval is smaller than this interval. Do you have some intuition for why these intervals are kind of asymmetric in this way? Why is it that the interval on the right is larger than the interval on the left? Yeah.

AUDIENCE: There's going to be certain kind of values that are right below 1/2 where with the bias in mind, our sender, they'd be more incentivized if the responder played high or believed that it was high.

And so for that kind of small range of values, that's why t1 is biased below 1/2.

PROFESSOR: Exactly right. Yeah. So if it was symmetric, then at the midpoint, the sender would always want to say the state is high. So we have to create these intervals to be asymmetric to ensure that the sender isn't tempted to say high.

And effectively what we're doing is we're shifting the midpoint left in order to counter that bias. So one way I like to maybe think about it is, this interval is smaller, which means the sender is giving more precise information about the state.

So in order to deter the sender from claiming the state is high, what we're saying is, well, if you say it's high, you're not able to convey very much information, and therefore the decision isn't going to be very precise. So the sender is trading off the bias of the decision against the informativeness of the decision. That's one way of thinking about it.

And then one final thing. Wait a second. t1 needs to be between 0 and 1. So we immediately see that beta has to be pretty small in order for this to make sense at all. Is that a question? Yeah.

AUDIENCE: Can you explain the whole thing about biasness about the decision, and what benefit there is from the decision in the lower part versus being biased in favor?

PROFESSOR: I guess that's-- I mean, I think that maybe another way of saying it is by having this interval larger, we're basically shifting the midpoint of this interval farther away from me.

So as the person right at the threshold, I'm tempted to claim the state as high just because I'm biased and I want the state to be higher. But as we enlarge this right interval, well, the decision that is induced when I say high starts getting higher and higher because the interval is so big and the midpoint is moving farther to the right.

And eventually, in the extreme case, the interval is going to be so big that I actually don't want to send the higher message. And the interval is going to be just the right length. That kind of balances those forces.

So maybe I guess the point that I'm making is the length of the interval determines how far the midpoint of the interval is from the endpoint of the interval.

And if we made that interval really, really big, then the midpoint of that interval is going to be really, really far away from its left point. And in fact, it's going to be too high even for the biased sender. So we have to balance these just right.

AUDIENCE: OK.

PROFESSOR: So I interpret the midpoint being far away from the point as saying, because I'm not revealing much information about the state, the receiver's updated belief about the state is very far from the true state. But maybe that's not so important.

So one observation here is that-- note that this only works if 2 beta is less than 1/2, which means beta is less than 1/4.

AUDIENCE: I have a question.

PROFESSOR: Yeah, absolutely. Yeah, questions are great. Yeah.

AUDIENCE: If t1 moves to the left, then the midpoint of the right interval goes down because it's 1 plus t1.

PROFESSOR: Right. So what I'm thinking about is the relative position of the midpoint relative to t1. So what happens is if I move t1 to the left, t1 moves to the left by, say, an inch. The midpoint moves to the left by only half an inch. So the difference between the endpoint and the midpoint grows.

AUDIENCE: And then at a certain point, because of the way the receiver is going to respond, it's a smaller difference from the left side?

PROFESSOR: Exactly. And that's exactly what we've done here. We found-- I mean, this formula is exactly telling us where that point has to be. So you're right, maybe I should have been more precise. It's true as I move t1 left that the high decision gets smaller.

But it doesn't get smaller fast enough, because when you move the endpoint by 1, you only move the average by 1/2. And that's exactly-- I mean, you can see that here. So I guess what I'm saying is y of H is t1 plus 1 over 2.

If we shrink t1-- if we reduce t1 by a little bit, then we reduce y of H by only half of that. But t1 gets reduced by the full amount.

AUDIENCE: And then so is the idea that the left side is smaller than the right side also just based on the idea that there's a beta?

PROFESSOR: Exactly, yeah. I mean, I guess very simply, all it's saying is we're choosing the intervals so that the midpoint of the left interval and the midpoint of the right interval-- if we look at those two midpoints, they average to the cutoff point plus beta. And that's exactly what we wrote here. I mean, this is the key condition. Yes.

AUDIENCE: Is the responder aware of the bias? And if so, are they incentivized change their behavior, given that the sender is going to be lying part of the time when they say high?

PROFESSOR: OK, great question. So first, this goes to an earlier question. They do know the bias exactly. And this is an equilibrium. So they recognize what the sender is doing.

But I wouldn't say the sender is lying. I mean, what the sender is doing is the sender is telling them whether or not the state is below t1 or above t1. And given that message, the receiver is behaving optimally.

Because if they get the message L, whatever the bias of the sender, they believe that the sender's message means the state is uniform over this interval. And therefore, their optimal decision is t1 over 2.

So I mean, you're thinking exactly about the right way. This is what the equilibrium is capturing. But the receiver doesn't have any incentive to deviate. Yeah.

AUDIENCE: That makes sense.

PROFESSOR: Yeah. Yes, other questions?

AUDIENCE: So does the receiver know t1?

PROFESSOR: Oh, good question. Yes. So t1 is a parameter of the equilibrium. So t1 is going to be known. And the reason is we-- or this gets into a subtle question of what's known.

But in equilibrium, when the equilibrium is played, the players have correct beliefs about what the other players are doing. So the receiver believes that the sender is following this equilibrium strategy for this particular value of t1.

So the receiver, in a sense, knows t1. Yeah. Yeah. So I guess we have to be careful. I was writing down all these formulas to try to derive the value of t1. But once I've solved for t1, formally, the equilibrium is this exact value of t1.

It's not an abstract parameter anymore. It has to be 1/2 minus 2 beta. Otherwise, the equilibrium wouldn't work. Yeah. Great.

So we said this only works if beta is less than 1/4. And indeed, you can show that if beta is more than 1/4, the only equilibrium is the bad no-communication equilibrium over here. So let's make this point over here.

If beta is greater than or equal to 1/4, no communication is the only equilibrium. Now, when I say only, I mean, the message m* that we use could be a different message, but we're never going to be able to convey information.

And the point here is the bias is so strong that it's impossible to credibly convey information. And you can think of this in the extreme case as kind of like the story of the worker who goes to the firm and says, I'm a good worker.

If the advisor has such a bias relative to the receiver and they always want higher decisions, the receiver just can't listen to anything they say. And therefore, the only equilibrium is going to be no communication.

If beta is less than 1/4, well, we can have some coarse communication where we reveal what the state is above or below this threshold. But we might ask, well, can we convey even more information? Could we have more precise information conveyed?

Do you have ideas how we might get beyond this case of just above or below a threshold and actually convey-- how might we convey even more information in equilibrium? Yeah. Yeah.

AUDIENCE: I had a question. So I was just wondering, for the utility functions, because we say that t is between 0 and 1 and then we also have this bias, is the sender always incentivized to-- like if t plus beta is ever greater than 1, wouldn't that mean their utility is always negative?

PROFESSOR: Yeah. So one thing is I wouldn't-- yeah, this is maybe a good point. You might worry-- the way we've written the game, the utilities are actually always negative.

Remember, with utilities, we can always add a constant to them and it doesn't really change things. So you might be thinking, oh they'd rather not play the game at all. They'd rather just not participate.

You could imagine the utility being 100 minus this and 100 minus this. And then the utilities would always be positive. So you can always-- the sender is only looking at differences in utilities, not the absolute levels.

So we could always shift up the utility and it's not going to make any difference. Yeah. So if you wanted, we could call this 100 minus this and 100 minus this, and it wouldn't change anything. Hopefully that answers that.

So any ideas? We have this equilibrium where I just say the state is either below a threshold or above a threshold. How could we maybe convey more precise information about the state? Yeah.

AUDIENCE: Still along the threshold idea, but is it possible to have low, medium, high?

PROFESSOR: I think that's a natural thing. Why don't you just say it's low, medium, high? Or maybe we could go a step further and split it up and say it's either really low or a little bit low or a little bit high or really high.

And it's not clear where this stops. We could do a lot of this. So that's going to be the idea. And these are called partition equilibria. Maybe I'll erase this.

So more generally, we're going to look at what are called partitional equilibria. This was a special case of a partitional equilibrium where we just had two cells.

We had two messages that said either below or above a threshold. But more generally, here's our interval 0, 1. Why don't we do something like this? We could say t1, t2. Maybe it keeps going. tk minus 1.

And we could say either the state is in this interval or it's in this interval or it's in this interval. It keeps going-- or this interval or this interval.

So basically our message by the sender is going to reveal which of these intervals the state is in. And as the number of intervals gets really big and each of the intervals gets really small, we're actually converging to more and more precise communication.

So let's understand what's going on here. It's kind of convenient to call this t0 and call this tk. So formally, I'm going to look at a partitional equilibrium with k cells.

So this is a k-cell partitional equilibrium where k is greater than or equal to 2. So we have the first cell, the second cell, the third cell, all the way up to the k-th cell.

And that means we're going to have k messages. And let's label them. Let's just call this message 1, message 2, all the way to message k. So what the sender is going to do is they're going to look at the state and they're going to say, well, if the state is between 0 and t1, I'm going to send message m1.

If it's between t1 and t2 I'm going to send message m2. If it's between t2 and t3, I'm going to send message m3, all the way up to mk. And maybe I'll call this y1.

If the receiver sees message m1, they're going to take decision y1. If they see m2, they're going to take y2, and so on. So formally, what I mean is yk is equal to y of m of k. But this is just cleaner notation.

And again, we're going to try to solve for these parameter values t1 up to tk minus 1, we're going to solve for these thresholds to see if we can make this an equilibrium. So what's going to have to happen? Well first, let's use receiver optimality.

Receiver optimality says, well, when the receiver gets this message m1, they know the state is somewhere between 0 and t1. And they believe, in fact, the state is uniformly distributed between 0 and t1.

So their updated belief about the expected state is that it's the midpoint of this. So in each case, the receiver's decision must be the midpoint of the cell that they're told the state is in.

So in math language, what we would say is yk or yj-- let's make sure we get this right-- is going to be equal to tj minus 1 plus tj over 2.

Because when they get message mj, they learn that the state is between tj minus 1 and tj. So maybe we'll say this here, because they learn that the state t is between tj minus 1 t and tj. OK?

So that gives us a formula for each of these decisions. This is for j equals 1 all the way up to k. And then what does sender optimality tell us?

Well, it's going to be the same idea. At each of these split points, we want the sender to be indifferent between sending the lower message and the higher message. So when the state is t1, the sender should be indifferent between m1 and m2.

When the state is t2, the sender should be indifferent between m2 and m3. All the way up when the state is tk minus 1, they should be indifferent between mk minus 1 and mk.

So what that means is that tj plus beta should be the midpoint of yj and yj plus 1 over 2. Let's make sure we got kind of the fence post problem right.

When the state is t1, we should be indifferent between m1 and m2, which is yj and yj plus 1. So that makes sense. And if we did a little more-- and this must be true for j equals 1 to k minus 1.

Notice we have k conditions for the receiver because we have k cells. We have k minus 1 conditions for the sender because we have k minus 1 split points. This is the classic kind of fence post issue, that k minus 1 cuts gives us k pieces.

So we could do a bit of algebra. I don't think we really need to go through all the algebra. Let me just kind of describe what you're going to find.

So maybe I'll just say mathematically what this is called is-- OK, we get kind of a system of equations for t1 through tk minus 1.

Remember, we already did the case of k equals 2. And in the case of k equals 2, we just solve for t1. But more generally, we have to solve for k minus 1 values.

But I want to just maybe convey the idea intuitively. Here's how I like to think about it. We could do the algebra, but let's just think about it.

Let's look at adjacent cells, OK? And let's look at what happens when the state is here. OK? Well, if these adjacent cells had the same length, would this-- let's call this tj.

Would the sender be indifferent between getting this point and this point? Which would they prefer? Well, these two decisions are equidistant from tj. But the sender is biased.

The sender prefers the decision that's closest to tj plus beta. So the sender would always prefer the higher decision. So what we need is we need tj plus beta to be equal to the average of these two points.

So what we need to do is we need the average of these two points to shift right by beta. Because if we want the sender to be indifferent between these two points, they shouldn't average to tj. They should average to tj plus beta.

So what we want to do is we want to move out this interval, make this interval longer, until the average has shifted right by beta. But we're only moving this point. This point is staying fixed. So if we want the average to go up by beta, we have to increase this point by 2 beta.

AUDIENCE: Why are we keeping the left point?

PROFESSOR: Say again.

AUDIENCE: Why is the left point fixed?

PROFESSOR: We're basically solving the equation forward. We're saying, let's suppose the left interval is like this. How big should the next interval be?

AUDIENCE: OK, because [INAUDIBLE].

PROFESSOR: So we have to increase this point by 2 beta to increase the average by beta. But how can we increase the midpoint of this interval by 2 beta? How much do we have to increase the right endpoint to get the midpoint to go up by 2 beta?

4 beta. So that means this has to go up by 4 beta. We could do the algebra, but I just kind of like to think it through. So what we indeed have is the general formula is this.

And you can solve it, but I think it's nicer to think of it like this. Let's write it this way. So we're comparing the lengths-- ah, no, no, that's wrong. Sorry, let me rewrite this.

So we get this formula, which compares the lengths of the consecutive intervals. So in our picture here, this is an interval length. This is the next interval length.

And what we're saying is, as we move to the right, the next interval moving to the right must be as long as the interval to the left plus 2 beta through the argument that we did over here. Yeah.

AUDIENCE: When you originally first drew the picture, that was saying like tj plus 1 minus tj was the same length as tj minus [INAUDIBLE] tj minus 1 [INAUDIBLE]. And then basically [INAUDIBLE].

PROFESSOR: Right. So formally in this picture, I fixed tj minus 1. I fixed tj. And I was saying, how big must tj plus 1 be? Maybe it would have been easier just to the algebra. But I find this clearer. Yeah.

OK. And actually, we can check. Let's go back to our example over here. When t1 was equal to 1/2 minus 2 beta, what were the lengths of these intervals?

This interval was 1/2 minus 2 beta. What was the other interval? Well, it's 1 minus that, right? 1 minus 1/2 minus 2 beta, which equals 1/2 plus 2 beta.

So indeed, that formula held in this special case. What that midpoint was doing is it was exactly ensuring that the length of this interval was exactly the length of this interval plus 4 beta.

And where did beta over 4 come in? Well, 4 beta has to be less than 1. And that's where it comes from. So now we can see what the structure of the equilibria have to look like if we you want to draw it more carefully.

As we move to the right, the intervals are going to keep getting bigger by a fixed amount. So it's going to be something like this. The first interval has some length delta.

The next interval has some length delta plus 4 beta. And the next interval has some length delta plus 8 beta. And we can keep going.

And then we have to make sure that everything works out at the right interval and everything comes together. So what's the takeaway? What we actually find-- let me just give the statement over here.

So it can be shown-- we won't go through it-- that there exists a partitional equilibrium with k cells if and only if beta is less than-- and I think this formula can really capture a lot of the key insights of the paper. So let's try to understand what's going on.

First, let's check that this is consistent with what we did before. So what if we put in k equals 2? With k equals 2, we're back to the case, we already looked at. And with k equals 2, the theorem says there exists a partitional equilibrium with two cells if and only if beta is less than 1 over 2 times 2 times 1, which is 1/4.

So indeed, it agrees with what we got over here. More generally, we get a formula like this. So what does it say? Well, as k increases, communication becomes more precise. As k increases, we're partitioning the state into smaller and smaller cells, and therefore conveying more and more precise information about the state.

So if we want more precise communication, what it's telling us is beta has to be even smaller. So that requires less bias.

So what we see is that the less bias that the agent is, the more precisely it's possible for the sender to communicate with that agent and more we can split the state space into these very, very tiny cells and convey more and more precise information. And in the limiting case, when there's no bias, we can basically perfectly reveal the state. Yes.

AUDIENCE: Can you explain a little bit more about the reasoning between precise communication coming from less bias or the last arrow?

PROFESSOR: Right. So I'm saying the last arrow comes from this equation. If k grows-- so we want to have an equilibrium with more cells-- then this inequality has to be satisfied.

Because k is on the bottom of this fraction, as k goes up, this fraction gets smaller. And therefore, beta has to be smaller in order for this equilibrium to work.

AUDIENCE: Is there a reason that's not related to this inequality?

PROFESSOR: Oh, yeah. I think the reason not related to the inequality is that when we provide really, really precise information about the state, it becomes very tempting for the sender to misreport and claim the state is higher than it is.

So think of the logic of that opening case where we fully reveal the state. If we basically say what the state is down to really, really strong precision, what's the sender going to do?

They're just going to claim the state is higher, they're going to claim it's in a higher cell, and they're going to benefit from that. If they're less biased, then it's feasible for us to split the state space into very fine cells, and the sender won't have an incentive to deviate.

So I guess one final picture I'll show. And I know this lecture is maybe a bit harder, so I think I'll end maybe 5 or 10 minutes early.

But what we can look at is I think it's helpful to look at the space of biases. So let's look at beta. And let's see what's possible in this game.

So here, we have 1/4. And if beta is up here, then no communication is possible. The agent is so biased that there's no equilibrium in which any information about the state is revealed. That's what we already said over there.

Let's get the numbers right. If we have 3, we have 6, it's already 1/12. Yeah, I think we'll get down to something like 1/12.

So here, there is a two-cell equilibrium only. If the bias is between 1/12 and 1/4, then the only equilibrium that can be shown is the kind of equilibrium we described, where you either say the state is above or below a threshold.

And we're going to keep going down and down and down. And then here in this case, we're going to have a two-cell and a three-cell equilibrium. So for a lower level of bias, there will be a two-cell equilibrium, where the sender just says the state's above or below a threshold.

There will also be a three-cell equilibria, where the sender says the state is either low, medium, or high. But that's it. There's nothing else.

And then if we get down to here, there's going to be a two-cell, a three-cell, and a four-cell. And as we keep going, as beta keeps getting smaller and smaller, the set of equilibria grows. We have more and more equilibria.

And we have equilibria with more and more cells that provide more and more precise communication. And the limit as beta gets really close to 0, we have very informative equilibria.

And maybe one final observation I'll make is let's suppose beta is small enough so that there's a range of equilibria. There's a one-cell, there's a two-cell, there's a three-cell, there's a four-cell.

The question is, which equilibrium is best for the players? And the intuition for this. Let's say we're in a case-- let's say beta is here. So we have both a two-cell and a three-cell. What do you think is better? Yeah. Or do you have a question? Yeah.

AUDIENCE: I was going to say, I think for the receiver, it's probably better to have more cells. So they're getting closer.

PROFESSOR: Exactly right. So it's definitely better for the receiver, because if there's more cells in the equilibrium, the receiver is just getting more information about the state. And therefore, they're able to make more accurate decisions.

It may be less obvious, but it's actually better for the sender as well. So it turns out the highest k is best for both the sender and the receiver.

So for any value of beta, depending on the value of beta, there's going to be a range of equilibria. There's going to be some uninformative equilibria. In fact, there's always the babbling equilibrium where you don't reveal anything.

There might be a two-cell where I just say it's high or low. There might be a three-cell where I say it's high, medium, or low. A four cell-- the biggest number of cells that can exist as in equilibrium, that's going to be the most informative equilibrium.

And it turns out that's going to be the equilibrium that's best for the sender and also for the receiver. So maybe that's some sort of prediction about what we might expect to see. Yes, Amy.

AUDIENCE: So it makes sense that it's best for the receiver just because when you get closer, the receiver [INAUDIBLE].

PROFESSOR: Yeah.

AUDIENCE: And how does the sender benefit? Is it also the same reason?

PROFESSOR: Yeah. It's actually a property of the quadratic law. So I think the intuition is that because the receiver is always matching the decision with their expectation of the state, the sender can never shift the decision.

So we're getting into a technical point. But it turns out the expected decision is always going to be equal to the expected state. And as the cells get smaller, the benefit is that the state is conveyed more precisely, and therefore higher decisions are taken in higher states and the lower decisions in lower states.

And that's going to be better for the sender. But the basic intuition is that the-- I guess the formal answer would be it's like a bias variance decomposition. You can express the expected loss as the bias of the error and the variance of the error.

And the bias is fixed and the variance you're going to make smaller. But I'm happy to talk more offline, if you want to go into that in more depth.

So I think that was a bit of a harder lecture, so let me stop there. I know we're near the end of the semester. I'll give people an early stop. And I will see everyone on Tuesday for the last lecture.