NANCY KANWISHER: OK, so we are now going to do the third part in our sort of trilogy of uniquely human functions. So starting a few weeks ago, we moved from things that are shared with animals, like navigation and basic visual perception-- very similar between humans and monkeys. And other number-- all of these things that are shared with animals, to talk about functions that are uniquely human.

And that includes music and language. And today, we're going to talk about the coolest and most quintessentially human mental function and that is thinking about each other's thoughts. So to get started thinking about this, let me remind you guys of what should be pretty obvious-- that we human beings are profoundly social in many, many different respects.

So if you think about-- if you recount what happened in your day to someone else, pretty much all of the elements of what you're recounting our interactions with other people. These are the things we care about. These are the basic structure of our lives. Other people are the source of our greatest happiness for most of us. If you ask people at the end of life what matters most, they'll say I shouldn't have wasted all that time working so hard.

And my job didn't matter. And I didn't realize, actually, it's just other people are the only thing that matter. Interactions with other people, or lack thereof, are also the source of one of the deepest kinds of human suffering for many of us, whether it's an everyday thing like a breakup with a lover, or a loss of a loved one, or a family member-- and also depriving people of social interactions.

This is, like, society's kind of strongest form of punishment, is solitary confinement. Interactions with other people-- our failure to understand them is devastating in the case of autism, where people struggle to understand what other people are all about. And in terms of our amazing abilities as human beings, most of what we know we learn from other people. Sometimes, we figure stuff out on our own.

But if you take an inventory of all the stuff you know, you read about it or somebody told it to you. You learned it from other people. And consequently, some of the greatest feats of humanity, from arts to science to everything else, are products of people working together. OK, further, many people have hypothesized that social cognition, or ability to understand and deal with-- interact with other people-- is one of the strongest drivers of the evolution of the human brain.

So many have argued that the hardest problem we solve on a daily basis is understanding other people. And it's also a very important problem. And so that's a powerful natural selection force, understanding other people. And that has shaped the structure of the human brain. Not everybody agrees with that kind of extreme view, that it's the primary driver, but it's an interesting hypothesis. And it surely played some role.

Social cognition also is just a very large percent of human cognition. And it's a large in terms of minutes of every day. If you tally up different forms of social cognition, this is one we're doing right now-- and we're about to engage in an hour and 20 minutes of it right here. And probably, when you leave this room, you'll engage in other kinds of social cognition. So it's just a large percent what we do all day every day, and it's also a large percent of what the cortex does.

So roughly that stuff and other bits are engaged in different aspects of social cognition. OK, so what exactly is social cognition? To get a sense of what is entailed in social cognition, which is a very complex, multi-part, multifaceted enterprise, I'm going to show you a short video of an interaction between a bunch of 18-month-old infants and an experimenter.

And what I want you to do is just watch it. It's very charming, but also think about what these kids must know to be able to do what you're about to see. Can you see the screen there? OK.

OK.

[VIDEO PLAYBACK]

That little kid in the corner doesn't know this guy. He's just been brought in for the experiment, and he's just watching this weirdo doing this thing.

Never met that guy before. Here's another case.

He's dropped a clothespin.

- Ah!

Oh.

Oh.

NANCY KANWISHER: What is the matter with this guy?

[LAUGHTER]

- Oh.

NANCY KANWISHER: Looking more and more suspicious.

[SPEAKING IN FOREIGN LANGUAGE]

OK, so what did you notice? What are those little kids, 18 months-- oh, there's another one. This is cute, but we don't need it. Let's go on. What did you notice about what those little kids were able to do? What is entailed in being able to do that? What have they figured out-- lots of things. Yeah--

AUDIENCE: They understand intention.

NANCY KANWISHER: Yeah. Yeah-- yeah, all they're doing is watching this guy do these clumsy actions. And he's not saying anything, he's just grunting. And they're figuring out what he intends, absolutely. What else? Yeah--

AUDIENCE: --sort of being able to interpret the sounds of grunts that the experimenter was making-- so especially, like, a lot of confusion. Like, the-- sound.

NANCY KANWISHER: Uh-huh. Yeah, one could ask which of those things are more important-- the actions or the sounds. I don't think they took it apart here, but it's an interesting question. Yeah so they understand some kind of vocal communication-- not language, but some kind of intent. That may be part of it. Yeah?

AUDIENCE: Yeah, I meant how to be helpful without doing-- something for real?

NANCY KANWISHER: I didn't hear. Without-- sorry.

AUDIENCE: Without interfering with whatever's coming in their way.

NANCY KANWISHER: Yeah, they know how to be helpful. But not just they know how-- they have a will to be helpful. Remember, these kids never met that stranger before this video started, or maybe a minute or two before. So it's fascinating-- not just what they're able to figure out, but what they're motivated to do. Right, that's a quintessentially human thing, is to spontaneously help a stranger who you've never met before and are not related to.

OK, all right, so-- but focusing more on the cognitive abilities. What you need to do to figure out another agent's actions is to figure out what that agent is doing. And that's externally observable. You can watch. You can use your perceptual system, construed broadly, to figure out what is this person doing in some general sense. And that's important, and subtle, and multifaceted, but it's nowhere near enough.

Really, what you want to know is-- why are they doing that thing? OK, and that's not directly observable. You can't see the why of a person's actions. You have to infer it. In fact, to infer it, you have to infer a bunch of hidden mental states. And those things are much more abstract than just, I'm picking up the phone. You know, that's just an action. You can see it.

But if you try to think about why am I picking up the phone, now we're in a whole different ballpark. And so what you need to figure out, to figure out the agent's hidden mental states-- you need to figure out what can that agent see or hear. OK, and here, it's really important to not lose a level here. We're not talking about what we can see. In understanding another agent, we need to understand what that agent can see about someone else.

OK, so we're getting kind of meta here-- multiple levels. OK, we also need to understand the other agent's desires and goals of what they want, or what they're intending, OK. So how might we figure these things out? Sometimes, there are simple cues that might suffice. So you might make an inference like, if a person is reaching for an object, they want that object. OK, so there, the externally observable thing might give you a pretty direct inference about the internal hidden state of what the agent wants.

OK, but a lot of times, that's not going to work. Right, that's going to work in some cases, but only a few. And further, the percepts and the goals that we've mentioned here aren't enough. OK, they'll work in some cases, but humans do much more than that. So let's consider this case. Here's an agent who's doing some action, and we can observe their body motions. And what we want to understand is how this other agent watching them is going to infer things about this guy.

OK, so we're interested in what's going on in the mind of this guy, as they watch that guy. OK, now let's consider the case that this guy is trying to understand-- this is Romeo here. Why did Romeo reach for the bottle? OK, well to understand why Romeo reached for the bottle, you need to perceive and infer a whole bunch of things. You need to see the hand reaching for the bottle to know what happened, but then to really understand the case of Romeo, you need to understand that his intention is to drink the liquid in the bottle.

That's why he reached for it. OK, you can't see the intention to drink the liquid. You can only see the reaching. Then, you further need to understand why is he intending to drink the liquid. Well, because he believes the liquid is poison.

I don't know why my pop ups are all messing up here-- whatever. He believes that Juliet is dead, and he wants to die. OK, so all of that is part of what is entailed in understanding why Romeo reached for the bottle. OK, and crucially the stuff in red, these things are-- the beliefs and desires are not directly observable. They're highly abstract. And they're crucial. They are the best way to explain and predict behavior.

So when we need to understand other people and what they're doing and why, we're in this abstract space of beliefs and desires. We can't see them directly. We have to observe. We have to infer them indirectly. And they're essential, because that's what we need to do to understand other people.

All right, so we talked before about perceiving an agent-- I'm sorry, about inferring what an agent can perceive, and inferring the agent's desires and goals. But as I just pointed out, we also need to know the agent's beliefs-- what they think. So all of this stuff in here, inferring the hidden mental states of agents is what we mean by mentalizing, inferring other people's mental states.

OK, in the cases we're talking about here, we're inferring the percepts of another agent, inferring the desires of another agent, and inferring the beliefs of another agent-- OK, all this abstract stuff. OK, no computer system can do this, not even close. See, we can still say that. We can't say that about object recognition anymore, but we have this little precious reserve where we can still say that.

No animal can do these things, except in very restricted cases. So there's an active, ongoing research program and a bunch of labs testing chimps in subtle tasks to try to figure out when they can infer things about the mind of another chimp or a person. And there are a few restricted cases where they can make some inferences about the mental contents of other agents, but the current dominant view is that those inferences are very restricted to very particular situations, usually involving competition over food.

And they don't generalize in any way. OK, so in this whole enterprise, specific cues like making an inference that reaching for x means wanting x will help in some cases, but will only get us so far. OK, we do much more than that, and the question is how. All right, so inferring mental states to understand other people involves inferring their percepts, their beliefs, and their desires.

OK, so how do we do this? Well, a first question that we ask in this class, and it's kind of a sensible thing to do, is we were trying to understand this whole space of how people think about another agent's percepts, beliefs, and desires. We might just think, OK, maybe that's just part of generic cognition. Who says that's a separate domain of cognition that we should study separately? Maybe we should just study thinking in general.

Thinking about objects, thinking about physics, thinking about people, maybe it's all the same thing. OK, so let's start by considering, is it really a separate thing from the rest of cognition? OK, well, there's a classic behavioral paradigm in this field that's provided some evidence in this. You may have heard about this in other classes, but we'll talk about it, because it's so kind of fundamental to all of the work on thinking about other people's thoughts.

And that is called the false belief paradigm. It's really a way of testing beliefs. And the reason you test false beliefs rather than true beliefs is if you ask about another agent's true beliefs, then your prediction about what they're going to do next is confounded with what is true in the world. And if you want to really tap into what's going on in their mind, as opposed to what would be true in the world, you have to use a false belief to pull those apart.

OK, so the classic way this is done, the Sally-Anne task goes like this. There are many variants of it, but this is the classic, original version that appears in hundreds of papers. So here's Sally and Anne, and Sally puts her ball in the basket. OK, you show this to little kids-- or animals, too, versions of it-- we're doing the little kid version. You say OK, here's Sally. She puts her ball in the basket.

I don't know why my pop ups are screwing up. Anyway, then Anne comes along. Sally leaves the room, Anne comes in, and moves Sally's ball from the basket to the box, and closes the top of the box. Then the question is, when Sally comes back, where will she look for her ball? OK, where will she look for her ball?

In the basket, right. OK, if you do this very simple question for three-year-olds, they'll say in the box. If you do it on five-year-olds, they'll say in the basket. OK, adults say in the basket-- you have to think for a second, right. If you just kind of blurt out the first thing that comes to mind, you'll behave like a three-year-old and say in the box. But if you think for half a second, you realize, duh, she doesn't know what got moved to the box. She'll look where she thinks it is-- in the basket.

OK, everybody with the program here? OK so that's a very, very simple task. It's the basic false belief paradigm. And to make it more vivid, I have to show you the video of Rebecca Saxe giving this task to a bunch of kids and talking about it. It's a delightful video, and I'll give you a really vivid sense of how smart a three-year-old is, and how much they understand, and yet how fundamentally and totally they fail this task-- versions of this task.

Come on, here we go.

[VIDEO PLAYBACK]

- The first thing I'm going to show you is a change between age three and five, as kids learn to understand that somebody else can have beliefs that are different from their own. So I'm going to show you a five-year-old who's getting a standard kind of puzzle that we call the false belief task.

This is the first pirate. His name is Ivan. Do you know what Pirates really like?

- What?

- Pirates really like cheese sandwiches.

- Cheese. I love cheese.

- Yeah, so Ivan has this cheese sandwich. And he says yum, yum, yum, yum, yum. I really love cheese sandwiches. And Ivan puts his sandwich over here on top of the pirate chest. And Ivan says, you know what? I need a drink with my lunch. So Ivan goes to get a drink. And while Ivan is away, the wind comes and it blows the sandwich down onto the grass. And now, here comes the other pirate. This pirate is called Joshua.

And Joshua also really loves cheese sandwiches, so Joshua has a cheese sandwich. And he says, yum, yum, yum, yum, yum. I love cheese sandwiches. And he puts his cheese sandwich over here, on top of the pirate chest.

- So that one is his--

- That one's Joshua's.

- And then--

- That's right.

- And this his is on the ground.

- Yeah, that's exactly right.

- So he won't know which one is his.

NANCY KANWISHER: Spontaneous! Don't even have to ask!

- Ivan comes back and he says, I want my cheese sandwich. So which one do you think Ivan's going to take.

- I think he's going to take that one.

- Yeah, you think he's going to take that one? All right, let's see.

- I told you.

- Oh yeah, you were right. You know--

So that's a five-year-old who clearly understands that other people can have false beliefs, and what the consequences are for their actions. Now I'm going to show you a three-year-old who got the same puzzle.

And Ivan says, I want my cheese sandwich. Which sandwich is he going to take? You think he's going to take that one? Let's see what happens. Let's see what he does. Here comes Ivan, he says, I want my cheese sandwich. And he takes this one. Oh-- why'd he take that one?

- His was on the grass.

- So the three-year-old does two things differently. First, he predicts Ivan will take the sandwich that's really his. And second, when he sees Ivan taking the sandwich where he left his, where we would say he's taking that one because he thinks it's his, the three-year-old comes up with another explanation. He's not taking his own sandwich, because he doesn't want to, because now it's dirty on the ground. So that's why he's taking the other sandwich.

Now, of course, development doesn't end at 5. And we can see the continuation of this process of learning to think about other people's thoughts by upping the ante and asking children now not for an action prediction, but for a moral judgment. So first, I'm going to show you the three-year-old again.

Was Ivan being mean and naughty for taking Joshua's sandwich?

- Yeah.

- Yeah.

- Should Ivan get in trouble for taking Joshua's sandwich?

- Yeah.

- So it's maybe not surprising-- he thinks it was mean of Ivan to take Joshua's sandwich, since he thinks Ivan only took Joshua's sandwich to avoid having to eat his own dirty sandwich. But now, I'm going to show you the five-year-old. Remember, the five-year-old completely understood why Ivan took Joshua's sandwich.

- Was Ivan being mean and naughty for taking Joshua's sandwich?

- Yeah.

- And so it's not until age 7 that we get what looks more like an adult response.

Should Ivan get in trouble for taking Joshua's sandwich?

- No, because the wind should get in trouble.

[LAUGHTER]

- The wind should get in trouble for switching his sandwiches.

[END PLAYBACK]

NANCY KANWISHER: OK, so everybody got the idea of the false belief tasks-- passed in the four and five-year-olds, not in the three-year-old. And the three-year-old not only fails to get the answer right, but when asked why it happened, he comes up with a totally different account. OK, and you can see that three-year-old is not a dummy. Three-year-olds are smart and can do all kinds of things, but they don't get this particular thing.

OK, all right, what about kids with autism? Well, kids with autism either failed this task altogether, or they pass it much later than neurotypical kids. OK, but now we've got to figure out why. Why would a relatively high functioning kid with autism-- say, a 7-year-old who's got language, and who seems to understand the question-- why would they fail this task? Why would anybody fail this task?

There's a couple of things going on here. The essential element of the false belief tasks that we're interested in is that it requires you to attribute thoughts to another agent. Right, that's what this task was made to tap. But you might fail it not just for failing to attribute thoughts correctly. You might fail it for other reasons. You might fail it because this task involves this weird situation where there's the other agent's belief, and there's a reality, and they're different.

And that's kind of confusing. It's like representing x and not x at the same time. Maybe that's just a particularly generically hard cognitive thing to do, OK. And the third possibility is that the true state of the world is so salient and dominant, it may be hard to just inhibit that true reality in order to infer the belief-- a belief is a kind of less visible, salient thing than the true state of the world. And so maybe just having an inability to inhibit dominant, salient representations would interfere with your ability to do this task.

So people with autism, or, in fact, three-year-olds, might fail for any of these reasons. So how are we going to figure out whether people with autism are actually failing this task because of their difficulty in attributing thoughts? Well, the way to figure that out is to come up with a really, really clever control task. This task was invented by Debbie Zaitchik when I was in graduate school. She was my office mate in graduate school, and she came up with this task. And I said, that is really brilliant.

So we want to figure out whether it's the belief, per se, that's the difficulty in attributing the thought or the belief to another agent that's the crux of the problem. So in Zaitchik's false photo task, the idea is it's a logically isomorphic task, but it's about a physical representation, not a mental one. So what you do is you show Sally putting her ball in the basket. And then you have what was then, like, widely known to kids, a Polaroid photograph-- a camera come along.

We didn't have cell phones back then, so it was a Polaroid camera-- takes a picture of Sally putting her ball in the basket, and then Anne moves the ball from the basket to the box. And then you ask the kid, where will Sally look for her ball-- I'm sorry, where is-- sorry, then you ask the kid, where is the ball in the photograph?

OK, sorry, I screwed it up. Here's a belief version. The kid watches Sally and Anne. In this case, you do a photograph version of it, so there's no-- you're not asking about another person's beliefs. You're using a photograph. And you're showing Sally put the ball in the basket. You take a photograph of it, Anne moves the ball to the box. And then you ask in the photograph, where's the ball?

OK, you're not asking about a mental belief, you're asking about a physical representation in the photograph. OK, so do you see how-- even though I bungled it on the first time-- it's really a logically isomorphic task. In both cases, you're asking about a representation that differs from reality, but one's a mental representation. That's where the kids with autism fail the task. This is the physical representation.

Kids with ASD fail the false belief task, not the false photo task. And that rules out these other accounts, because in both cases, there's a representation that differs from reality. You have the x and not x challenge-- right, that's true of both of these. And in both cases, the true state is kind of more dominant than the representational state in the belief, or the photo.

And yet, they have a problem with the belief version, not with the photograph version, suggesting that this is the correct account of why they fail. That make sense? Was there a question? Yeah-- OK. So all of that is consistent with the idea that attributing thoughts to another agent is a special, distinctive thing that can be-- that's what I just said.

So we have evidence that attributing thoughts is a distinct domain of cognition. It's just not part of your generic ability to think about anything. And we see that from the fact that typical children have this very systematic developmental time course, where they get this at age five, they don't get it at age three. And we see it from the selective loss of this ability in kids with autism, that they get the false belief task way later than they get the false photograph task.

Was there a question here?

AUDIENCE: So the kids are never asked what does Anne perceive. That question then comes up in the FP task-- in the false photograph task?

NANCY KANWISHER: I'm sure people have done versions of that task, but tell me what you're getting at here.

AUDIENCE: I just want to know whether-- to do an accurate comparison, I would assume that the final question remains the same. So I would ask the kid, what do you think Anne perceives given the change in the stimulus. And then the kid responds, given the new material, as to what the answer--

NANCY KANWISHER: You're trying to get a tighter control to the false photo task.

AUDIENCE: Very much, yeah.

NANCY KANWISHER: Yeah, instead of where will she look for the ball-- yeah. But it's not really about what she perceives. It's about where she thinks the ball is.

AUDIENCE: Sure.

NANCY KANWISHER: And so the false belief tasks have been done every which way, from where does Sally think the ball is, to where will she look for the ball. And you get the same answer in all of those. So I think that's closer to, where is it in the-- where is it in the photograph is really, where does she think it is? That's the parallel. Yeah, because she's not still perceiving it at that point.

OK, so that's behavioral evidence to suggest that there's really something different about inferring another person's thoughts and beliefs than there is from the rest of cognition-- the distinctive developmental time course and this kind of selective deficit in autism.

So what can we learn from functional MRI? Is there a special part of the brain for making these inferences about other agents, thoughts, and beliefs. OK, so Rebecca and I did this experiment 100 years ago where we scanned people while they were doing simple verbal tasks. So we wrote false belief stories, and sort of false photo stories-- generalized-- and we scanned people while they lay in the scanner and just read these simple descriptions and answered basic questions.

So you can try this right now. Here's a false belief story. You read-- Susie parked her sports car in the driveway in the middle of the night. Nathan moved her car into the garage to make room for his minivan. Susie woke up early in the morning. Susie expects to see in the driveway-- a sports car or a minivan? Yes.

Sports car, duh. Right, exactly. What she thinks is there, not what is there. The control experiment that-- we wrote many, many of these-- a volcano erupted on this Caribbean island three months ago. Barren lava rock is all that remains. Satellite photos show the island as it was before the eruption. In the photos, the island is covered in rock or vegetation, right.

OK, so here again, it's just like Debbie Zaitchik's ASD control here. It's-- you're asking about a representation in the satellite photo that's no longer true of reality. And you have to distinguish the two. So we scan people just doing these tasks. OK, and if you look in a whole brain group analysis that I was much dissing before, but that's a good first step, you see-- here's a slice through the brain like this. You see a bunch of hot spots that respond more when you do these tasks than those tasks.

OK, and there's a bunch of these regions. Deep down in the middle of the prefrontal cortex here that you can see in a slice like that here is a region called MPFC for medial prefrontal cortex. It does all kinds of interesting things, and it will get mostly short shrift in this lecture. It activates in this task-- for reasons you'll see in a minute, activates in this task. Oh god, pop ups have a mind of their own. But here's another region, the right TPJ-- TPJ stands for temporoparietal junction. So in me, if my temporal lobe is going down here, my parietal lobe is here, frontal lobe is there.

Temporoparietal junction is somewhere right around there. OK-- everybody oriented? And if we look at that region, shown symbolically here, what you see in the time course of response in that region doing this task, is here's what happens in that region. This is time here-- many seconds, as they read and answer that task. And here's what happens when they read and answer the false photo task-- so false belief task, false photo task. That just shows you the same thing's activation. You can see it evolve over time.

OK, so what happens in the MPFC region-- oh god, why is this not happening.

Is that showing up? Oh, there it is, OK. Right-- in the MPFC region, you see a similar thing. You get a response below baseline here. We'll talk about that-- what it means to have a response below baseline in some conditions in the next lecture. It's actually not all that mysterious. It just means that baseline, remember, is usually lying in the scanner staring at a dot-- going dum dee dum dee dum-- or whatever you do when you're staring at a dot in the scanner.

And the key point for now is just that you also see a higher response to the belief than photo case. OK, so what's going on in these regions? This activation, this higher response, to the belief than photo case-- even though I was touting what a brilliant control condition the false photo task is-- but it's still open to a bunch of different possibilities. OK, so those activations could be just thinking about a person.

After all, our false belief stories all involve people, and our false photo stories did not. So maybe this isn't about beliefs per se, it's just about thinking about people in general-- any aspect of people. OK, or it could be any mental inference we make about the internal, invisible states of a person. OK-- anything that might be going on in their head. Or it might be, particularly, attributing thoughts and desires to the person.

OK, so how are we going to unconfound these-- with a new experiment and some new conditions. So here's a new condition that refers just to a person, not about their internal invisible states, or their thoughts or desires. Subjects just read simple stories like-- I see what's happening. It's not showing-- it's showing something different on the screen from what's on my slide here. I have no idea why. But anyway, that's why I keep wondering what's there and what's not.

OK, so in the external case, in this condition, you lie in the scanner and you read stuff, like, Andrew just had a growth spurt, so he was gangly and rather awkward. Like most teenagers, he had bad skin and bad taste in clothes. He wore mostly baggy jeans and flannel shirts. OK, so it's not riveting, but it's not boring either. You read a bunch of these, they're interesting enough. OK-- but no mental states there, just outward appearance.

OK the next condition refers directly to the person's thoughts-- Nicky knew that his sister's flight from San Francisco was delayed 10 hours. Only one flight was delayed so much that night, so when he got to the airport, he knew what flight was hers. OK, not even that interesting, but it's talking about what he inferred, what he knew-- beliefs. OK, but there's a third condition that refers to an internal state, but not a thought or a belief. It refers to a visceral state.

Sheila skipped breakfast because she was late for the train to her mother's. By the time she got off the train, she was starving. Her stomach was rumbling and she could smell food everywhere. OK, so this case is interesting. It's vivid. It refers to a mental state, but not a thought. It's a physical state, hunger. OK, I mean it's a mental representation of a bodily sensation. OK, everybody get the distinction between internal thoughts and feelings, sensations?

OK, so this contrast is consistent with any of these hypotheses. And we're going to test those hypotheses with these three conditions. OK, so here are the three hypotheses we just described here, and here are the three conditions I just gave. And I want you guys to tell me what we're going to predict that we will see in the right TPJ according to each hypothesis-- namely, which of these conditions will produce a high response, and which will produce a low response.

OK, so let's start with this case-- on the hypothesis that the right TPJ responds whenever you're reading about a person, what are we going to find? Which of these conditions is going to produce a high response?

AUDIENCE: All of them?

NANCY KANWISHER: All of them, yeah, absolutely. So that's the prediction here-- all of them. OK, what about if it's interested in internal bodily sensations, or actually, in any internal mental states?

The last two, absolutely-- both the visceral case and the thoughts case. Sorry-- that cell and that cell. OK, and what about if it's specifically about attributing thoughts to another agent. Just the last one, right-- OK, very good. So what do we see? So first of all, here's just a localizer task. Always nice to look at your localizer and make sure it worked out. So this is a response of the right TPJ to the belief and photo condition, just like I showed you before-- higher when you do the belief task than the photo task.

OK, just a reality check-- key question, what is the response of the right TPJ in that main experiment with the three conditions we just described? And here's what it does. It's high only in the thought condition, only when you're reading about, thinking about another person's thoughts. And that's significantly higher than either the external properties or the visceral states. And those two, even though it looks like there's a little bit of a trend that visceral is higher than external, that's not significant.

Yeah?

AUDIENCE: On the previous slide, is it, like, thoughts and/or desires?

So like, why can't-- like, I was just confused about the example for the visceral case, if your smelling food everywhere. Like, couldn't that fall under it--

NANCY KANWISHER: That's a good point. Yes, it could. Yes, you could infer a desire in that case, absolutely. Yeah, that's right. So yeah, that's a good point. It's intended to just identify the vivid, visceral state. And it's probably hard to do that with it without invoking some kind of desires here. Yeah, I would consider it a confound for that, for that particular example. I'm guessing it was not so for the others. In fact, it sort of couldn't have been, given the way this came out, right? But good point.

OK, so--

OK, so-- and if you do a whole group analysis, the whole brain, and you say what bits do you get in a contrast of reasoning about another person's thoughts versus their external appearance, or their visceral states, their bodily sensations? You get the right TPJ again. OK, so everybody got this? This suggests that the right TPJ is extremely specifically interested in inferring another person's thoughts, not even just their bodily sensations.

So it's not all mental states of another person that engage this region. I mean, I find this quite remarkable, because that's so specific. Before you hear about this, you think, OK, thinking about another person's thirst or hunger or pain, is that really different than thinking about their beliefs? But oh yes, it's different. Right? The TPJ does believe it's not thirst and hunger and pain. Yes?

AUDIENCE: What about your own thoughts and--

NANCY KANWISHER: Yeah, so there's a whole literature on this. And I think if you-- OK, so first of all, we need to get the levels straight. When you're just thinking, you're not thinking about thinking. You're just thinking, right, so you're not engaging this region. So when we're talking about thoughts, we're talking about you thinking about another person's thoughts. So the parallel would be thinking about your own thoughts, right. So I think-- I can't tell you exactly what the literature is on this, I'm sure there's a few experiments.

But I think what you'd have to do is adopt that kind of meta perspective on yourself, right-- a lot of our thinking is sort of thinking about thinking, but not in a very explicit, separable way. I'm not being very clear on this, but I think if you ask people for example, when you saw this surprising event, whatever it is-- you make something up-- what were you thinking? You might engage that region. I realize I'm not totally sure if there's literature on this. Heather, is there? There must be literature on this.

AUDIENCE: [INAUDIBLE]

Yeah.

AUDIENCE: [INAUDIBLE]

You get MPFC for that, yeah. Or go ahead--

AUDIENCE: [INAUDIBLE]

AUDIENCE: --and very, it's kind of, like, intricate-- [INAUDIBLE] --it will two networks that are--

NANCY KANWISHER: But not just rTPJ, right, rTPJ and other regions. Yeah, so you could get some rTPJ when answering those kinds of questions.

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Uh-huh.

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: OK, so you get something in the vicinity. I'm just trying to get--

AUDIENCE: [INAUDIBLE]

Right, but the other part of this is, I bet you have to do something quite explicit to do that, because in some sense, we're thinking about our own thoughts kind of implicitly a lot of the time. Right, so I'm guessing that it's when you ask about it explicitly.

AUDIENCE: Yeah, I forget the details of-- [INAUDIBLE]

NANCY KANWISHER: Right.

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Right.

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Right. OK, anyway, good question. I'm sorry, I don't have a totally adequate answer yet. I'm guessing the literature doesn't have a totally adequate answer yet. OK, all right, so recall that I mentioned that the contrast between doing the false belief task versus the false photo task gets not just the rTPJ, it also gets this medial frontal region. OK, and that medial frontal region responds more to the belief than photograph task, like this. What does it do in this split between thoughts, visceral thinking about other people's thoughts, their visceral states, and their external appearance?

It doesn't care. OK, so there's a real division between these different brain regions that are engaged when you do the false belief task. The rTPJ is very specifically interested in-- very specifically engaged when you think about another person's thoughts, not when you think about their bodily sensations. The medial prefrontal region is engaged in all three of these conditions. OK, so what it looks like is that of these three hypotheses, any time you think about another person, that's true of the MPFC.

Whereas the right TPJ responds specifically when you attribute thoughts and desires to another person. OK?

OK, now you may have been wondering about the fact that all of these experiments are using words. It's kind of not, like, normal social cognition. You're reading about all this stuff. And if these regions were really doing what we're claiming they're doing, we should be able to find them in other situations where you make inferences about other agents' beliefs. So in more recent experiments, they've been showing Pixar movies to subjects in the scanner.

So you can just watch yourself make a few inferences about the agent's beliefs. Imagine you're watching this in the scanner, and just let's just do a little bit of it to get the gist.

[VIDEO PLAYBACK]

[MUSIC PLAYING]

No words, it's a key.

OK, so there's a whole little microcosm of mental states here. In other parts of this same six minute Pixar movie, there's very vivid bodily sensations. There's a porcupine who inflicts real pain with porcupine quills, so there's bodily states, and inferences about bodily states, and inferences about mental states that go on.

And by showing this movie to subjects in the scanner, and then labeling which parts of the movie require the viewer to make inferences about the thoughts and beliefs of the protagonists, and which require them to make inferences about bodily sensations-- like, usually pain-- if you do that contrast, you find the same set of regions, the rTPJ and its left hemisphere bit, and you find some medial prefrontal stuff, and you find some other stuff.

OK, so what does that mean? That means that you don't need words to identify this region. If you induce the same cognitive processes just from watching a wordless movie, you show the same selective activations. And that's cool, because it tells us that this is not about some kind of verbal reasoning. It's really about the kind of deeper cognitive process, whether you do it based on a movie, or based on a bunch of sentences.

And that's a powerful generalization. We like imaging results, where the result generalizes across very different kinds of tasks. And this really strengthens the evidence that this is really what's going on in this region. And it also means you can look for this region in kids. OK, and I didn't manage to fit that into this lecture, but there's a whole other research enterprise where, in Rebecca's lab, they've been scanning kids watching this Pixar movie, and asking how that region develops.

And they find that, actually, the region continues to develop even after age four. It continues to develop. And what happens is it doesn't so much get bigger as it gets more selective. In the younger kids you get activation, both for-- well, it's just less selective. Anyway, so this shows generalization. So now we've shown that we see the rTPJ is selective for thinking about other people's thoughts.

We see that with false photo-- false belief versus false photo contrast. We see that it's highly specific. It's not just any thoughts you have about another person. And we see it generalizes to Pixar movies.

OK, so let's consider now moral reasoning as a test case for theory of mind. Why moral reasoning? Well, if you think about it, reasoning about what's a morally acceptable or unacceptable action on the part of another person is all about what the person intended and what they knew. OK, so intent is very fundamental. It's built into the legal system.

Think about the difference between murder and manslaughter. Right, they both involve killing another person, but one is with intent, and the other is accidental. And our common moral reasoning and our legal system cares deeply about that difference.

So for example, I'm going to give you a moral reasoning task, and just think about this. And your task is going to be to decide how morally permissible is the action described here-- Grace's action in particular. OK, so Grace and her friend are taking a tour of a chemical plant. And when Grace goes over to the coffee machine to pour some coffee, Grace's friend asks for some sugar in hers. And there's white powder in a container by the coffee. The plot thickens.

The white powder is a very toxic substance left behind by a scientist, and therefore, it's deadly when ingested in any form-- versus-- but the container is labeled sugar, so Grace believes that the white powder by the coffee is sugar left by the kitchen staff. So grace puts the substance in her friend's coffee, and her friend drinks the coffee and dies. Now, your question is how morally permissible was Grace's action, on a scale from one, to totally not OK, morally forbidden, to seven, morally permissible.

OK, so think about that on a scale, and write down your number on a piece of paper. You don't have to divulge it. OK, everybody got the question?

Everybody decided, more or less? OK, write down your number. Now, consider a slightly different case-- slightly, but crucially different case-- this case is known as the accidental harm case. Now consider the case where instead-- OK, instead of being labeled sugar, you get the same story, but now the container is labeled toxic. So Grace believes that the white powder is a toxic substance left behind by a scientist. Nonetheless, she puts the substance in her friend's coffee and their friend drinks a coffee and dies.

Now, consider how morally permissible is Grace's action, from one, to totally not OK, to seven, morally permissible. And write down your number.

OK, how many people gave a lower number-- that is, more morally forbidden for the second one than the first one. OK, if you didn't, you probably weren't paying attention. OK, so you can see that it's the crux of the matter what Grace believed when she did the action. OK, that's why we're talking about moral reasoning here as a test case of theory of mind, because what the agent knew at the time of the action is of the essence in thinking about the moral status of their action.

OK, everybody got that-- OK, so this is a powerful test case. And notice that-- also in the clip that I showed you from Rebecca's Ted Talk, she showed that kids' ability to use an agent's knowledge in doing what you guys just did, that kicks in a little bit later than the standard false belief task. So sometime after they get the basic idea of false belief, they start to apply. It takes a while to kick in in this other case.

OK, so this is known as intentional harm as opposed to accidental harm. OK-- it's just the terminology in the field. OK, so what do you think will happen in autism? If we ask people with autism these two questions-- and what do you think will happen if we apply TMS to the right TPJ? It's right out there on the lateral surface, just asking for it. So it's a totally doable experiment. It's been done, what do you think happens?

OK, let's take the case with autism. How do you think people with autism will respond to these two questions? The same as you guys did? Yeah-- yes, Shardoul.

AUDIENCE: Probably the average distance between-- the average distance between how bad they think the second one would be less than the first one--

NANCY KANWISHER: Yes, exactly. Why?

AUDIENCE: Because they might not be able to make the distinction between Grace knowing that the container was toxic versus the container being labeled toxic. Like, both then, would be confounded with just like truth and consequence.

NANCY KANWISHER: Absolutely, absolutely. Did everybody get that-- share that intuition? To the extent that autism is a particular deficit in understanding what another person knows or believes, that's the only difference between these two cases. To the extent that you have difficulty representing that, you will have less of a difference in your moral judgment about these two things, because you have a hard time representing that person's knowledge.

It's not that autism is a deficit in moral reasoning. It's that moral reasoning entails thinking about other-- at least these cases, not all of them-- but these cases involve thinking about other people's thoughts and taking them into account. And to the extent that that's difficult for you, you will make less of a distinction, exactly-- right. Yeah, was there a question back there?

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Yeah, OK. OK, what do you think will happen if you zap the rTPJ with TMS while subjects are doing these tasks? Yeah?

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Yeah?

AUDIENCE: They'd be, like-- [INAUDIBLE]

NANCY KANWISHER: Yeah, yeah. That's the prediction. If the rTPJ is the main bit that's doing the inference about-- it's representing the beliefs of others, then if you zap it, you might change people's moral inferences-- moral judgments. That's pretty wild. That's what happens. Zap the rTPJ, and they make a smaller distinction between accidental harm and intentional harm. OK, so both of those things are true. I won't drag you through all the details of the experiments.

But the basic findings from this whole line of work show that, first of all, neurotypical people agree, as you guys did, that accidental harm is more morally permissible than intentional harm. OK, and people with autism give less forgiveness for accidental harm compared to intentional harm than neurotypicals. OK-- just because that ability to represent the key knowledge that tells you it's accidental is something they're not good at.

What is the role of the rTPJ? The data show that forgiveness for accidental harms-- first of all, I left this out before, is correlated in neurotypicals with activation of the right TPJ during moral judgment. So if you just measure across a whole bunch of those moral judgment problems how strongly activated was your rTPJ as you read that problem, that's correlated with your ability to forgive somebody for accidentally harming someone, again showing that there's a relationship between your representing the thoughts and beliefs of another person and your using that information to exonerate them from harm they didn't intend.

Yeah?

AUDIENCE: For someone with autism syndrome, are they-- is it different for thoughts versus external actions of the things?

NANCY KANWISHER: Yeah, so first of all, I'm treating autism very superficially here. It's an extremely heterogeneous thing that varies not just along a spectrum, you know, which it clearly does, but probably along many spectra, and is highly heterogeneous. So these experiments are just done in high functioning adults who are-- totally pass false belief tasks-- they pass them later in life, but they get to-- otherwise, you can't test them on these kinds of experiments.

And the effects are quite subtle. They're just a slightly lesser difference between accidental and intentional harm. OK, so just to clarify that, which I probably should have said. But your question is, are the deficits in autism specific to thinking about thoughts rather than thinking about actions?

This is ongoing work, but a lot of research has shown that it is more specific to thoughts, and that a lot of the stuff you read about just basic perceptual difficulties-- I mean, most studies find that people with autism are a little bit worse at face recognition, but not much worse. In tasks asking about goals of actions, like reaching for objects-- what is that person's intention? Mostly, they don't find a deficit in autism.

So the perceptual-- basic seeing people and seeing what they're doing is much less impaired. That's what the current literature suggests. There's always the worry that we're not asking in the right way, or testing in the right way. And the literature is highly inconsistent from one study to the next. I used to work on autism and I just couldn't stand it anymore, because every time a study's done, it gets the opposite of the previous study. I think as a population-- so heterogeneous.

But from a gloss, it looks like there's more of a deficit in the inferences about thoughts than inferences about actions. Yeah?

AUDIENCE: I just want to clarify. So the autistic subjects are also explicitly told that Grace believes-- blah, blah, blah. Yes? It's explicit--

NANCY KANWISHER: No, it's exactly what I gave here. Yeah, no Grace believes-- yeah, you're right. Grace believes. Yes--

AUDIENCE: Explicitly being-- and in spite of that, it's-- [INAUDIBLE]

NANCY KANWISHER: Yep, yeah, yeah. And so I think you can think of this as a subtle case. These are people who pass the explicit false belief task. But like the 7-year-old kids, there's one thing if you're asked-- what does this person believe, and another if you're asked a moral reasoning task for which you have to realize to bring the belief into account. It takes more of your own kind of active-- was there another question?

AUDIENCE: What if you ask-- like, instead of Grace, what if you say it's you?

NANCY KANWISHER: What if what?

AUDIENCE: You ask them about themselves. So you and your friends are taking a tour of the chemical plants, and give them the exact same scenario. Would that be any different?

NANCY KANWISHER: Yeah, that's interesting. It's like the question about to what-- sorry--

AUDIENCE: What was the question?

NANCY KANWISHER: Question is, suppose you ask people with autism this same question, but it's not about Grace. It's about you. You do all of this.

I'm not sure. Good question. Probably, somebody has looked at that. Yeah.

AUDIENCE: I'm not sure if you answered already, but do people with autism have something different in their rTPJ?

NANCY KANWISHER: We're getting there. You should absolutely be wondering that, and I haven't yet. And we're getting there. Good that you ask. In fact, that's probably my very next slide.

OK, so causal role-- we showed the causal role of the rTPJ in neurotypical subjects, you zap the rTPJ and you slightly reduce the difference in moral permissibility of the accidental harms and the intentional harms. OK, so all of these findings suggest that the rTPJ is causally engaged in understanding the difference between intentional and accidental actions, and that ability is specifically disrupted and autism, All. Of which leads to the natural prediction that Gisella just made about the rTPJ in autism.

So what do you think? Is the rTPJ affected in autism? How many-- raise your hand if you think there's going to be something different in the rTPJ autism versus typical subjects. OK, raise your hand if you think there isn't. Not sure-- OK, well, you're sort of both right in different ways. OK, so the answer is this-- so in Rebecca Saxe's lab, they did a study with a really large number of typical subjects. It's-- because of this heterogeneity in the autism population, it's really hard to get a stable result you can believe.

And it helps to have a really large-- it's hard to get enough autism subjects, so we try and we get as many as we can. And the most you ever get in a study around here is 20, 30, and that's a big struggle. But it helps to have a really large neurotypical control population to reduce your error bars on what the neurotypical population shows.

So this study is probably the biggest that's been done. They had 31 high functioning people with autism and 462 neurotypical individuals. They didn't just go run 400 new subjects for this. When you run that localizer task, you've got it in every study you run. And so then you can take all those localizer tasks across hundreds of subjects-- so-- and so what they find is, both if you look at regions of interest analysis, like find rTPJ, and if you do a whole brain group analysis, you find no differences between people with autism-- high functioning people with autism-- and typical subjects in the size, location, or response magnitude of the rTPJ.

OK-- when people do theory of mind tasks. So you should be surprised. Everyone is surprised. You all made the prediction. Everyone else made the prediction too. And that's pretty bizarre.

But does that mean that the rTPJ is not affected in ASD?

Shardoul?

AUDIENCE: If we feed it-- the information from the rTPJ is not used by other parts of the brain to-- [INAUDIBLE]

NANCY KANWISHER: That's a great hypothesis. Maybe the information is in there, and it's some kind of a disconnection thing, and so it can't be accessed by other processes. Right, absolutely. What's another hypothesis?

Yeah, Isabelle?

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Sorry?

AUDIENCE: It's being processed in a different order?

NANCY KANWISHER: In a different order. What do you mean?

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Yeah, but then we'd have to think about how the different temporal order would lead to the different behavioral outcomes. Right, it might. Yeah, David.

AUDIENCE: Maybe it's given a different priority, so, like, while you're processing there, I'm not being-- [INAUDIBLE] --the person itself [INAUDIBLE].

NANCY KANWISHER: Yeah, OK, so it might be there, but less salient or less important. But how is that going to account for the lack of a difference between the ASDs and typicals. Remember, you run your basic false belief versus false photo contrast, and you find the rTPJ. And surprisingly, even in this fairly large sample, it looks the same in size, location, and response magnitude in the ASDs and in the typicals.

Now, I left something out. This is high functioning adult ASDs who can now pass the false belief task. Right, that's crucial, because that's the task you're doing in the scanner. If they don't understand the task, there's no point scanning them kind of going, huh, what? Right, so these are people who are very high functioning. And they can totally do the task-- right, otherwise you can't run the experiment.

So-- yes?

AUDIENCE: What are the major differences that you're finding then, if it informs [INAUDIBLE]

NANCY KANWISHER: Well, that's a very good question. That's another reason that I stopped working on autism.

[LAUGHTER]

There's a whole battery that you run to try to establish that these people, officially, really, for the purposes of scientific study, count as having autism. And these people, really, officially don't. At MIT, you need to run those studies on everyone, because some of your control subjects end up in the other group.

And so those are a whole kind of battery of things, from involving an hour long interview with a trained person who tallies things like how much eye contact is made, and what kind of give and take happens in conversation, and all that kind of stuff.

AUDIENCE: I just don't-- I guess for the experimenters, I mean--

NANCY KANWISHER: Yeah.

AUDIENCE: --team. Their behavioral response might not necessarily be the rTPJ, but some other factor.

NANCY KANWISHER: Yeah, absolutely. So yes, it's possible that the rTPJ is absolutely fine in people with autism, and there's no difference, as those initial results seem to suggest, and that whatever differences you see with autism reside elsewhere. But that's surprising, given all the stuff I've said over the last hour about how, at least, in the case of moral reasoning, these same high functioning-- the same people who show the same activation of the rTPJ have slightly different moral reasoning, right.

So there's-- yeah?

AUDIENCE: My point isn't that. It's-- or so, for moral reasoning, for example, could be-- what they base their morals off of might not be determined by some other path-- I don't know. They might have a slightly different idea of what morals--

NANCY KANWISHER: Yep, that's true. Yep-- Nava?

AUDIENCE: This test, they functioned the same way-- the people who are being examined?

NANCY KANWISHER: Yes.

AUDIENCE: --in the same way. But in a different test, they performed different. So why don't they test them when performing those?

NANCY KANWISHER: Good, good-- exactly. So we're in this funny position of-- we've sort of identified theory of mind inferences as a critical difficulty in autism. But these are people-- the only ones we can scan on that kind of task are people who are going to already pass it. So we're already in a weird situation. Now, you might have predicted that they could pass it based on other cognitive abilities-- right, that they come up with another strategy.

And in that case, they wouldn't be using their rTPJ, but another hypothesis is it just develops later. They've got it. There it is. They use it. It activates the same. That's yeah-- and so your point is, why don't they test them on the moral reasoning task? Right, and see if it comes online? Right?

AUDIENCE: Well, the one that they can't function on, because--

NANCY KANWISHER: Exactly.

AUDIENCE: --they don't even know when to use it?

NANCY KANWISHER: That's right. That's right. That's right. It's a very good suggestion. That's probably been done. And I don't know this. Heather, do you know if people scanned ASDs doing the moral reasoning task? Probably Leanne, has done that right?

AUDIENCE: [INAUDIBLE] I don't know.

NANCY KANWISHER: It's a really good suggestion, and I actually now can't think why they wouldn't have done that. Do you mind just going on PubMed and look-- it would be Leanne Young. And it would be-- yeah, moral reasoning and fMRI. I bet she's done that. Maybe Heather will get you this, because it's a very good suggestion. Yeah, why not test them or the thing that's different.

OK, anyway, what I'm getting at here is-- all these hypotheses, you guys are raising are very good ones. But there's another one, which is, maybe the rTPJ isn't functioning right even though we see it activated more when they do the false belief task than the false photo task. Maybe we'd see a difference if we looked at the pattern of responses in there, right?

All right, so in this study, Koster-Hale and Saxe did something that should be very familiar to you guys. They took these two cases, sort of a version of what you're suggesting-- yeah, maybe it does entail it. I think of it as an MVPA experiment, but actually, Heather, that's the thing to look at-- does Koster-Hale see a difference in overall magnitude? And that I don't know.

Anyway, they did the MVPA version, so you have subjects do those same tasks-- the accidental harm, the intentional harm. You split your data in half. And you ask whether the rTPJ represents the difference between accidental harm and intentional harm. OK, everybody get why this is a sensible thing to do?

OK, so first you do that in neurotypical subjects, OK? And you find-- here is the-- remember, this is the original Haxby version, the correlation within versus correlation between. So this is a correlation within, and that's a correlation between. So this is accidental-- the pattern in the rTPJ for accidental harm, to accidental harm across stories, and intentional to intentional is higher than from accidental to intentional here.

I'm sort of skipping over the details, hoping you guys remember this. Is this making sense? OK, so typical classification with correlations, and you see that that was significant in this group of subjects, of typical subjects. OK, so that's cool-- so there's information in the rTPJ about whether an action was intentional or accidental. That's cool. Now we've learned something more than just it activates when you do these tasks. We know something about what it represents.

But you might say, oof, that's so teeny. Ick-- yeah, it was significant, but really? So what do you do in that case.

AUDIENCE: Do it again?

NANCY KANWISHER: Do it again. Absolutely. You don't go find some fancier stat that gets your p level, so you have to start-- no, no, no. You do it again. Rebecca and her lab members, being good scientists, did it again. And they got the same thing. Excellent-- new bunch of subjects, new bunch of stories, replicate and generalize, do it again. And so yes, indeed, we can-- the spatial pattern contains information.

Further, you can then-- again, in this bunch of neurotypical subjects-- look at the degree to which subjects rated the moral permissibility to be different in the intentional and accidental. Look at their behavioral ratings during the task. That's here. And you can see, that's correlated with the degree of pattern information in the rTPJ. OK, everybody get that? So the more you pay attention to that distinction, behaviorally, when you're doing the task, the more you think accidental harm is really much more OK than intentional harm.

The greater that difference in your behavioral ratings, the greater the discrimination ability in your rTPJ while you're doing it. OK, so that's a further link that's where the action is-- just a correlation, not causation. But it's a nice one. Yeah, OK-- so all of this is.

Yeah? Yes?

AUDIENCE: It's the [INAUDIBLE] there's no [INAUDIBLE] and there is a [INAUDIBLE] --so neurotypicals have a difference in [INAUDIBLE]

NANCY KANWISHER: Just univariate overall magnitude, accidental versus intentional, in neurotypicals, but not in ASDs. Aha, so it's not just a pattern result.

AUDIENCE: Interaction-- it's in the text. It's not in--

NANCY KANWISHER: OK, OK, OK. OK, but that's the--

AUDIENCE: --interaction of, not a-- it's just--

NANCY KANWISHER: OK, but then the version of Nava's question, which I resonate to very much, is if you just did the original kind of Dufour thing of-- is the rTPJ just as big and selective, and is it in the same place? Just--

AUDIENCE: If you did the Dufour thing, though, it [INAUDIBLE] there's a huge [INAUDIBLE] so you're unable to--

NANCY KANWISHER: That's true. OK, so they don't have the data. OK.

AUDIENCE: There's no way. You have to have the--

OK, fair enough. OK, good point. OK, so that's why-- you asked that question, it's a very good question. I'm thinking, yes, they should have done that. Reason they didn't do it is-- they could have the big sample because they were using localizer data, which they had from study after study after study. Every time you run a theory of mind test, you run that localizer. And so then you just go back a few years, and you've got 300 subjects, right.

But they didn't run the moral reasoning task on hundreds of subjects, and so they don't have the power to be able to see it. Yeah-- OK, fair enough. OK, so back to this. This is just showing that in the neurotypical subjects, there's pattern information in the rTPJ about intentional versus accidental harm. Everybody got that-- and it's correlated with your behavioral reading of the moral permissibility. OK, so that's cool.

So does rTPJ distinguish between accidental and intentional harm in ASDs? OK, so these are the data I just showed you. Actually, they did three experiments showing this here-- into neurotypicals. And here's the ASD data-- no difference at all. OK, so what that means is even though ASDs have the same size and location and magnitude of response of the rTPJ in the standard localizer task as typical subjects, the key difference that's been found so far is that in neurotypical subjects, the rTPJ holds information about the distinction between intentional and accidental harm.

And in ASDs, it doesn't. OK, that's probably just one of a bunch of things that are going to be different. This is, as far as I know, the only one that's been published. I see some squints. Are you not getting this, or do you have a question about it?

AUDIENCE: What happens to the other deficit in the behavior, sort of being called [INAUDIBLE]?

NANCY KANWISHER: Yeah?

AUDIENCE: Also, I have to ask, is this exactly the same?

NANCY KANWISHER: The false-- there isn't an obvious discrimination that you can do. See, the nice thing about the moral reasoning task is it's got these two outcomes-- these two conditions, accidental harm and intentional harm. So that gives us a way to go in and do the pattern analysis questionnaire. So with false beliefs, you'd have to think of some other dimension to look at.

And the experiments are not set up that way with stimuli that are on either side of a dichotomy, so that you can do the discrimination. You could do representational similarity analysis, actually, come to think of it. Then, you wouldn't have to have the whole dichotomy. That's interesting. OK, Heather, has Rebecca done that? You're my informant here. You could take--

OK, so Shash is suggesting, why not ask this kind of question of this standard localizer experiment? And I said, well, it wasn't set up with an obvious dichotomy. And then I was thinking, actually, you could take all of the belief conditions, and you could do an RSA on those in the ASDs and on the typicals, and ask if their patterns are different. Right? I bet they've done that. See what I mean?

Anyway, all right, I'll stop speculating and just ask. She's right there, two floors up. I can just ask. All right-- anyway, does everybody get this basic idea here? So the thing-- the rTPJ is there in the high functioning people with autism, but it doesn't hold the same information. OK-- all right.

All right, so where do we get with all of this? We used all of this stuff on moral reasoning as a way to look at the rTPJ and theory of mind. We found that people with ASD put less weight on a person's beliefs when judging the moral permissibility of an action. TMS to the rTPJ disrupts moral judgment. Pattern analysis shows that the TPJ distinguishes between intentional and accidental harm in neurotypicals, but not in people with ASD.

OK, so there's a nice little story developing here. I'm sure it's not the whole deal.