

[MUSIC PLAYING]

SARAH Today on the podcast, making science more open and inclusive and *so. much. better.*

HANSEN:

REBECCA SAXE: There are brilliant scientists all over the world who don't have access to an MRI machine and don't have the kind of money that it would require to collect those data. By making them publicly and freely available, scientists anywhere in the world who have a hypothesis can use our data to test the hypothesis and participate in science, even though they can't afford to collect their own data.

SARAH I'm Sarah Hansen. My guest today is the instructor for MIT course 9.401, Tools For Robust Science.

HANSEN:

REBECCA SAXE: I'm Rebecca Saxe. I'm a professor in the Department of Brain and Cognitive Sciences, and I'm Associate Dean of Science.

SARAH Rebecca is a pioneer in the world of science, not only in her social and developmental cognitive neuroscience research, but also for her advocacy in making science better, more inclusive, more open, more transparent. In this conversation, she digs into what it looks like to really push science to be more replicable and usable, how scientists can leverage data and contribute to fields outside their own for the sake of public progress, and we also talk about the power of wallowing in the classroom. I learned so much from her about the way science is conducted and how people like Rebecca and her lab are pushing the field forward. I hope you enjoy our conversation.

You're the principal investigator at the Saxe Lab. As far as I can understand, it's a lab that explores how humans think. I'm wondering if you could elaborate on that and tell us a little bit about what happens in your lab.

REBECCA SAXE: That's nice. Well, the research questions that interest me are about, as you said, just the big picture questions, how humans think, how we understand the world. I'm especially interested in when the problems we're solving are abstract. So we have to figure things out that we can't see or directly experience. And the kind of example of that is other people. How other people work, how they think, what makes them do their actions, how we're related to them. That is super important to absolutely everything in our life, but not directly visible.

So compared to lots of cognitive science that studies how do you see or how do you hear, which is all about processing the information that's coming in through your eyes or through your ears, like processing the structure of the visible or sensory world, I'm interested in processing the structure of the world that's invisible.

SARAH Wow. And how did you get interested in how we know or think we know how other people are thinking? Where did that come from in your life?

REBECCA SAXE: I'll start at the beginning. well, I think the instinct to understand the world we see in terms of the invisible parts that we don't see, I think I've always wanted to do that. When I first found out about science, that science was a thing, I thought maybe I wanted to be a chemist. Because the coolest thing I could imagine is that the physical objects in the world are made of molecules. That's just such a cool idea.

And then I thought maybe I wanted to be a biologist, because our bodies come from DNA. And isn't it amazing that these tiny things you can't see, the DNA, make the whole thing you can see? And then I found out about how thought comes from neurons. And I was like, that's the coolest one. The fact that the experience of our minds, all the thoughts we ever have, that comes from the pattern of fire-- the spatial and temporal firing of cells in our brains. Just that is the coolest thought ever.

So since then, I've wanted to study that, and I went to college to study that, how the mind works in terms of the brain. But that's a lot of different problems that you could study anything. You could study attention, you could study memory, you could study language. And all of those things are incredibly interesting. But yeah, I guess I got lucky. I got interested in specifically social reasoning.

And why partly it feels incredibly important, it's part of our life all the time. Partly, as I said, I'm most interested in the problems our brains solve without being able to directly experience the evidence. And partly the reason I do this work is a coincidence or maybe luck. So I was a PhD student here at MIT. I came to MIT in 2000 to join the lab of Nancy Kanwisher. She is a vision scientist. She studies how the brain processes information that we can see.

But when I was her graduate student, she gave me a chance to study whatever I wanted. She's an incredible mentor. And she thought we should study how you see people move, the visual part of social life, seeing people's actions. And we did start trying that. But I wanted to also try studying the invisible part, how you think about the mind.

And then what happened was we did both kinds of experiments. We did some experiments that were about seeing social movements and some experiments that were about thinking about invisible minds. And the thing is that the ones about the invisible mind work better. It turns out there is just this really strong signal in the human brain whenever people are thinking about other people's thoughts. And it was so strong and so reliable and we could see it in every single person that we tested. And so almost against our will, all of the papers we published were about those experiments, because those were the ones that worked. And so that started my whole career.

SARAH So how do you visualize the invisible? What do you see that tells you that a person is thinking about another person in their social environment?

REBECCA SAXE: You mean how do we see inside somebody's brain?

SARAH Yeah, how do you know as a scientist in a study that part of their brain is activated and thinking about what another person is thinking about?

REBECCA SAXE: Yeah, so there's two parts of that. One is, how do I know what part of your brain is active? And the other part is, how do I know why it's active? So how do we know that it's active? The main tool that we used was fMRI, Functional Magnetic Resonance Imaging. That's an incredible tool that lets us measure blood oxygenation as it changes over time, as you think different kinds of thoughts. And it was, again, an incredible stroke of luck for me that when I started grad school, that tool was new. And so many things about the human brain were unknown. Up until the late '90s, almost all studies of the human brain that we could do about what different brain regions did, well, there were a few tools available, but most of our knowledge came from brain damage. You would wait for some brain damage to occur.

And so our knowledge depended on tragedy. And also, fortunately, brain damage is relatively rare. And also the selective kind of brain damage that's most informative that happens only at one part of the brain region at a time and in young people, all of that is rare, especially in very young people. So if you want to study development, fortunately, there's not that much focal brain damage in young kids. That means we didn't know that much about how different brain regions acquire their functions in humans. And so we relied on studies of animals.

Well, fMRI transformed what was possible. And fMRI is a safe and pretty easy way to measure brain activity as people do whatever you want them to do. And the early fMRI studies, again, were about basic sensory functions. Looking at images, hearing sounds, tapping your finger. But it immediately became obvious you can do almost anything inside an MRI machine. So we could study any aspect of thought.

So we use fMRI to study brain activation based on blood oxygen. And initially in what we called typical human adults, who were MIT undergraduates. And then later in a broader array, including children. And now in my lab, we're studying infants. So we're using fMRI to study the origins of these brain regions in human babies. So that's one half is how do we measure brain activity.

Your second question was how do we know why that brain region is active? We can see blood oxygen. How do we know why that brain region is active right now? And that's experimental design. So over many years of experiments, we had people do different kinds of tasks inside an MRI machine.

So sometimes we would write stories that had different content and we would track, OK, this brain region, if you read a continuous story, it's only active when the story starts to be about the mind of the character. If the story is about the character's body or physical environment or what they're doing or what they're wearing, then that brain region activity stays low. But if we start to talk about what they thought or what they wanted or how they felt, now that brain region's blood oxygen goes up. So that's one way.

Another way is we would have people watch movies and we could do the same thing. We could look over the course of the movie, when are the moments in the movie that are really about the character's mind? And we could watch that the brain region's activity goes up at those moments. So for example, we have a big data set that I'm very proud of, which is adults and children watching Pixar's movie *Partly Cloudy*. I don't know if you know this one.

SARAH I don't. I'm going to have to watch it tonight with my 7-year-old.

HANSEN:

REBECCA SAXE: You definitely should. It's six minutes long.

SARAH OK, done.

HANSEN:

REBECCA SAXE: Yeah. And wordless. So the central drama of *Partly Cloudy* is about a stork in a cloud. You may know clouds make babies, storks deliver them.

SARAH Sure.

HANSEN:

REBECCA SAXE: So the core plot of *Partly Cloudy* is about the relationship between one cloud and one stork. The cloud, Gus, makes dangerous babies. And so the stork alternately has physical experiences of being rammed by this baby ram or electrocuted by a baby electric eel or poked by a baby porcupine.

SARAH Oh gosh.

HANSEN:

REBECCA SAXE: So very physical experiences. But there's also a drama about emotions, about betrayal and trust. And so we could look at how the movie alternates between focus on the body and focus on the mind and see that correspondingly the brain regions that we had identified are correspondingly going up and down inside that movie.

SARAH Wow, that is so fascinating.

HANSEN:

REBECCA SAXE: Thanks.

SARAH And so what's the difference between mind and thought? And is there a difference?

HANSEN:

REBECCA SAXE: OK, I mean, I guess how do we use those words in English? I would say the mind, that's a noun that refers to the collection of thoughts. A thought, I guess, is a piece. You'd say one particular thought has a particular content. I'm thinking this conversation is interesting. That thought is one element. And one mind is like one person, One mind is me. You're another mind. Whereas I have many thoughts. Does that help?

SARAH Yeah, it is interesting. It seems fundamental to trying to understand other people. They're so complex.

HANSEN:

REBECCA SAXE: Yes. And so you're right that part of what we study is thinking about minds. That is, thinking about different people and their different minds. And part of what we study is thinking about thoughts in the sense of inside you are many thoughts, and I might need to work out what each of those are to figure out what you're going to do next.

SARAH Yeah, and that process in and of itself feels like such a scientific playground to try to figure out how we know

HANSEN: those individual thoughts within another person's mind. Like, wow, that could be a lifetime of work.

REBECCA SAXE: It's a good problem, yeah.

SARAH Yeah. I'm curious if your lab is interested in the effects of screen time on children, like spending more and more

HANSEN: time on screens and less time interacting with people has on brain development.

REBECCA SAXE: Yeah, here's a funny study from my lab or a story. So in 2017, a postdoc came to my lab and we were trying to think what would be an interesting research program for her to do? She had a couple of years to be in my lab. And based on a study we had read in mice by a colleague, we thought, you know what would be interesting to study is the brain effects of social isolation. What happens when you're forced to be alone?

And so we did this study. We had all these people come into the lab. And one day we had them go without food, and one day we had them go without social contact. And then we measured their brain activities. And we--

SARAH Sorry, these are mice or these are people?

HANSEN:

REBECCA SAXE: We studied people. Other people had studied mice.

SARAH So you didn't feed them for a day?

HANSEN:

REBECCA SAXE: That's right. So we had people who either--

SARAH I'm assuming IRB signed off on this.

HANSEN:

REBECCA SAXE: Absolutely. We had people who went without food for a day or they went without social contact for a day. And what we were asking is we know that if you go without food for a day, at the end of that day, you'll be craving food. But if you go without social contact for a day, at the end of the day, will you be craving social contact? And is that the same brain region? Is there some brain region in your brain which is craving a basic need? That will crave food if you're hungry, or will crave people if you've been socially deprived.

SARAH Oh, so interesting.

HANSEN:

REBECCA SAXE: So that was the question. Thanks. Yeah, so this study took years to do, because we had to have these people come to the lab. We actually had them come three times, once with no deprivation, once hungry, once socially isolated, and we had to scan them. And yeah, it was an elaborate paradigm in order to scan them without social contact and so forth. Anyway, it took years.

We eventually had all the data. We analyzed them. We found out actually our hypothesis was right. There is a shared brain region that responds when craving and that responds to food when you're hungry and people if you're socially isolated.

SARAH It's the same region?

HANSEN:

REBECCA SAXE: Same brain region. Yeah. I mean, there's lots that's different, but there is one underlying brain region that's the same.

SARAH Wow.

HANSEN:

REBECCA SAXE: And it's very similar to in mice. So this social need is possibly a conserved feature of being a social animal. Anyway, the point of this story is we got that result in January of 2020. And I remember saying to my postdoc, the main problem with this paper is I don't know what it's about. Why will anybody care what the brain effects are of acute social isolation? That's not a thing.

People study loneliness, chronic isolation. But somebody who has had a rich social life and suddenly for exogenous reasons, has to be alone against their will, that's not a thing. How are we going to tell anybody why they should care about our results? And in March 2020, I was like, you know, I don't think we have that problem anymore.

SARAH Wow.

HANSEN:

REBECCA SAXE: That really did happen.

SARAH Wow. Oh my gosh. So let's talk a little bit about your commitment to open science, especially as that plays out in your lab. And then in a little bit, we'll get into the course. But you've made an explicit commitment to supporting open science through promoting transparency and reproducibility. Can you tell us why that commitment is important and what it means in practice in your lab?

REBECCA SAXE: Yeah, those are great big questions. And they connect to the biggest and deepest ideals of science for me. So to be a scientist, I think what we're trying to do is build cumulative understanding, where if I do an experiment or make a discovery, I can pass on to you what I did, what I observed, and what I inferred from that in a way that you can take the next step. You don't have to do my step again. You can get something from what I did and what I learned and build on it in many possible ways.

You might build on it by limiting it, by saying that's an interesting discovery, but it only happens in this context, or there's limits on its generality. Or you might build on it by asking a different question using a similar method. You studied that this way. I'm interested in this other thing. There's all kinds of ways we want science to be cumulative, but most of them require that I can give you what I did and my confidence in some way that you can find useful.

OK, so that's this big aspiration. I think both the way that we make claims about truth, like what we mean by scientific claims aspire towards truth. I think that's related to being open to certain kinds of scrutiny or open to scrutiny in different ways and what we mean. So to me, OK, scientists want to make discoveries that are true and useful in some degree. Somewhat true and somewhat useful, maybe more true or more useful. And truth, I think, is very much related to being open to scrutiny and useful as being open to reuse.

So for me, the commitments to transparency and openness in science are just related to the deepest notions of what it is to be a scientist and to do science. That's the abstract, big picture, pie in the sky. Now down to the nitty gritty.

SARAH Yeah, what does it look like in the lab?

HANSEN:

REBECCA SAXE: So in practice, historically the way that scientists communicated about their science was mostly through papers. So I would do an experiment, and then I would describe what I did in words. And the way that you could figure out what I did is from the words I described. So I would describe my method and you would read those words. I would describe my observations of the data and you would read those words. And that was how I passed on to you what I did and what I found.

And both because of disadvantages of the limited amount we could convey through words and the biases that were introduced by those limits, but also because new technology afford new opportunities, it's not necessary anymore to only communicate that way. So if you want to know about my experimental paradigm, there's so many more ways I could share it with you than just describing it in words in a PDF and mailing you that PDF.

I could make a video and put it on YouTube or in JOVE, the *Journal of Visualized Experiments*. I could make my entire experimental paradigm be code, which I put on GitHub, and then you could rerun it. I could build my paradigm in an experimental website like Prolific or Qualtrics or Lookit for kids. And then instead of describing my experiment, I could just give it to you. Here's the experiment. This is the experience that a person in my experiment had. Here it is.

And that is such a more direct way to be open to scrutiny. If there's something wrong with my experiment, you will really be able to see it. You don't rely on me having happened to mention it. And also facilitates reuse, because now here it is. You could use it yourself or you could redo it. So that's for experimental paradigms.

The same is true of data, It used to be that we would just describe the results of our experiments. But now since those data are almost always digital, they could be posted in a repository. And then instead of only whatever tests I thought to do of my data, if the full data are available, you could test whatever you wanted. Again, that means it's open to scrutiny. You could check if I'm right. You could test alternative confounds in my analyses. You could learn from my analyses how to analyze your own data in some way. Or also maybe there's things in my data I never thought of.

And so I was telling you, for example, about this data set of children watching this Pixar short. So a wonderful, amazing, brilliant PhD student in my lab, Hilary Richardson, who now runs her own lab in Edinburgh, she for three years in her PHD, had kids aged 3 to 12 in an MRI machine watching the same Pixar movie. So by the end of that, she had 122 kids who had all watched the same movie but varied in age.

At that point, this is a number of years ago now, that data set is incredible. There's so many things that you could test in a data set of 122 kids watching the same movie while we record activity. We record simultaneously in every part of their brain. And Hilary and I had a particular question we were interested in about thinking about other minds. But that's not the only thing in these data. Every part of their brain's activity was recorded.

And so after we had tested our particular question, there were so many other questions you could ask in those data. So here's the choice scientists face and Hilary faced this. Here's this data set. It took years and many thousands of dollars to acquire. Here's the choice.

We could keep those data for ourselves and when we thought of a question, analyze it and write a paper and slowly over many years, anything we happen to think of, we could test in our data and publish papers on it. That is, in some sense, the traditional way science worked, We collected the data, we own the data, we make all the claims from the data. And some people argued even at the time that that would be in Hilary's interest. Her professional interest would be in hoarding the data she had collected in order to get out of it as many papers as she could.

And she, I think, very bravely, thought there's more scientific good that could come out of these data than I can do by myself. And some of it is on questions that are not my expertise, like vision and all these other domains. And so instead of keeping the data private, she made the data public. And they've been reused dozens of times. They've been downloaded thousands of times. They're in textbooks. They're in training programs for undergraduates learning how to analyze fMRI data use these data. Literally this morning I saw on social media a new paper published with a new method that used those same data.

SARAH It's so exciting.

HANSEN:

REBECCA SAXE: It is. And I just think there's so much potential scientific upside from giving new minds access to ask new questions. And it's obviously delightful that with their new method, they confirmed our claims about how the brain develops. That's great. But they also discovered things we would never have thought to ask.

And the other big difference it makes making these data public is, as I said, those data took years and thousands and thousands of dollars to collect. There are brilliant scientists all over the world who don't have access to an MRI machine and don't have the kind of money that it would require to collect those data. By making them publicly and freely available, scientists anywhere in the world who have a hypothesis can use our data to test their hypothesis and participate in science, even though they can't afford to collect their own data.

SARAH I'm curious if you could articulate for us the structures in place that are making it not so common that this

HANSEN: happens all of the time. Because the way you're describing it, it sounds like, well, all data should be open and all science that's at least funded publicly should be open.

REBECCA SAXE: That's increasingly happening and also required to happen. So as of a year ago, the Nelson memo from the Office of Science and Technology in the White House has mandated that all publicly funded data do need to be openly shared. So I think we're in an era where open sharing of data is both happening much more now, more appreciated in terms of its scientific value, and likely to accelerate.

The challenge is not will data be posted. They probably will now that it's required. The challenge is how to make them useful. Because a data set takes a lot of interpreting. From bits to meaning is a lot of work. And the work that goes into making a data file understandable to somebody else so that someone else could use it is a lot of work above and beyond the work of acquiring the data.

And so the current challenge that scientists are facing is not are we going to make our data public, because we are. The challenge is where will the infrastructure and skills come from to ensure that when the data are shared, they are findable and usable?

SARAH I feel like this is a whole new arm of what we mean by science communication that we need to teach

HANSEN: undergraduate and graduate students. I feel like now the emphasis is on how do you communicate your findings through papers, through presentations, through the new ways that you're describing. But I feel like going into the future, we also need to teach people how to use metadata, how to make the data in their raw form understandable by other scientists. It feels like that's going to be critically important.

REBECCA SAXE: I love that and I could not agree more.

SARAH I'm going to read something to you that you wrote.

HANSEN:

REBECCA SAXE: Oh dear.

SARAH It's very good. On your website for the Saxe Lab. You said, "Serious inquiry into human nature is not a luxury. No

HANSEN: one should be excluded from it by their race, gender, or other accidents of birth." I found that really important. And so I'm wondering how open not only applies to data practices, but also to community practices in your lab, perhaps ways you invite everybody from the community to participate in learning about mind or participate in studies or in the doing of science.

REBECCA SAXE: Yeah, I love that you are seeing a connection between open data and open science as a practice for scientists and inclusion and diversity. I think those are deeply related in the sense that, again, as I said, I think when we claim to be making scientific progress, part of what we're saying is these discoveries and ideas, they're good in a way that withstands certain kinds of scrutiny or they've been subject to certain kinds of scrutiny.

And diversity and inclusion is central to being sure that our ideas have been scrutinized in enough ways. When we are walled off in tiny, homogeneous groups, we don't get the scrutiny we need to find out if our ideas are rigorous and robust. And it's very easy for ideas to be accepted without being good ideas when they're never subject to scrutiny from all possible angles. So I do think there's a deep ethical connection between who we include in our scientific communities and what we mean by openness about our scientific products.

And so you said, how does that connect to my lab's practices? I hope it does. Aspirationally, it does. Of course, not perfectly and maybe not up to our ideals every day or in every way. But we do try to think about both separately and together who we include in our scientific conversations, how we are open to scrutiny about our scientific claims, and what those two ideas have to do with one another.

SARAH So you teach 9.401, Robust Science, which directly takes up the topics of open, transparent, and reproducible
HANSEN: science. So could you tell us a little bit about who takes this course and what you hope they will learn by taking it? The big picture learning outcomes.

REBECCA SAXE: Absolutely. I can tell you multiple versions of the answers to both of those questions, because there are many versions of the answer to those questions. So the original impetus in this class was between 5 and 10 years ago. As I was looking at scientific practice as it was then in my field, I saw a lot that was not ideal.

I saw a lot of practices that were standard and were viewed as necessary evils. Hoarding data or non-reproducible methods or-- people would say we want better scientific communication, but there's no incentive to do better scientific communication. Or similarly, diversity, inclusion, something we say we want, but that there was very poor progress on.

And specifically in my own fields, looking at cognitive science and cognitive neuroscience, so many papers that didn't have open data, didn't have open methods, had not been preregistered, and that I thought the rigor of the evidence really suffered for the lack of those practices. At a certain point, it was incredibly depressing to use all of my hard earned expertise to carefully read a very technical paper and upon careful reflection decide I had no idea if any part of it was true. Because the way that science was communicated just didn't let us rigorously evaluate it. And there was a reproducibility crisis happening where I knew from my own practice.

So I should briefly say the class I taught before this was an undergraduate lab class in which we tried to reproduce recently published findings in the psychological literature. So I knew from experience that a large number of recently published findings could not be replicated. And among other things, I also knew that many fMRI studies could not be replicated, that the ways in which we were incentivized to publish our claims were not well aligned with rigorous scrutiny.

SARAH It must have been like, what are we even doing here?

HANSEN:

REBECCA SAXE: It was incredibly depressing, to be honest. Because as I said, many of us become scientists with big ideals about truth and cumulative science and making something useful. And to realize that we were engaged in this activity that produced many papers but not a lot of reusable findings, it was really depressing and it led to a deep existential crisis.

There was a phase when I had students in my lab, I would say, come to lab meeting with a paper that, based on reading it, you're sure would replicate. That's what we were looking for. And we would just read paper after paper and be like, no, I don't think that would replicate.

So anyway, we got ourselves very depressed. And we were reading about the fact that science is not diverse and inclusive, the fact that data are not shared, the fact that findings are not preregistered, and therefore have all these biases in them. And I just got more and more depressed and nihilist.

And then I thought, I can't do this and I can't teach this way. I can't teach my students or my lab when we leave the room depressed every time. I can't do that anymore. So I resolved to make a class whose form would be every week, we would take a topic. We would give serious due consideration to what was wrong to the ways in which our practices did not live up to our aspirations.

And then we would switch and look at tools we could use to make it better. And we would practice using those tools and we would consider how well they worked and we would consider what was still needed from those tools to make them better so that every class would end in a positive mood.

SARAH So this is not only science, this is self too.

HANSEN:

REBECCA SAXE: At some point recently I thought, apparently the way I treat my existential crises is through teaching graduate seminars. [LAUGHS] I feel like that might be the most academic thing ever.

SARAH I think so, yes. But it's working.

HANSEN:

REBECCA SAXE: In response to existential need, I teach graduate classes. But you know what? It's so great. It has totally worked.

SARAH Yes. And it's helping science.

HANSEN:

REBECCA SAXE: It might be helping science. And it's really helping me. Because every week, we take on a problem and we wallow for the first part of the class.

SARAH Good, I love wallowing.

HANSEN:

REBECCA SAXE: Yeah, we talk about how bad it is. And then in the second part of the class, we talk about all the things people are doing to try to make it better and our experience of using those tools. And so we leave every week with new tools and new practices and new ideas for how to make our science better.

SARAH That's so great. Can we talk about one or two examples of what you wallow over and a strategy that students

HANSEN: might leave with? For example, the literature is biased.

REBECCA SAXE: Yeah, I thought you were going to go there. Yeah, so that is a big one. So in many types of science, and this is definitely true in the kind of science I do. So we'll talk about fMRI again. That's the main tool that I used for the first 15 years of my career. So in that tool over time, like over two hours, I measure change in blood flow in every part of your brain simultaneously. So I end up with a huge data set, hundreds of thousands of points in your brain measured over two hours of time. It's many, many, many, many data points.

And then I have some idea about where and when change will happen in your brain. But there are many ways to analyze those data. There's so many choices. There's choices about how what's called pre-process. So, for example, parts of the signal that I measure are actual brain activity and parts of it are noise. Noise because you moved your head, noise because you were breathing, noise because you clenched your eyelids.

There's all kinds of things affecting-- noise because the machine was heating up, for example. So many things that I measure are not related to your brain activity. So I have a whole bunch of choices about how I separate the data into the part I think is real, brain, and the part I think is noise. That's a bunch of choices. Many choices.

Then let's say we have the part I think is brain. There's again many choices about how I model that activity. There's choices about what shape do I think the blood response will have over time. There's choices about am I measuring individual brain regions or am I measuring their interaction, and what term am I going to use for their interaction? There's choices about am I measuring the amount of activity or the pattern of activity?

And if this gives you a sense that there's many choices, I'm vastly understating it. There's thousands of choices that we make before we give an answer where we say this brain region was more likely to be active whenever the story had a person in it, for example. So we are making those thousands of choices.

The problem is, say I have an idea of what I think the brain activation should look like, and I make one set of choices for how I'm going to analyze the data. And I look at the data and I don't find evidence for my hypothesis. I might think, well, that's probably that I just made a wrong choice. Maybe I needed to look at the interaction between the brain regions, not just one brain region at a time.

Or maybe, this is a real example that happened, I was trying to find a particular brain region and I used this strategy for finding it. I think it didn't work. I think we missed it. That's why the data aren't confirming our hypothesis. We didn't actually find the brain region we were looking for. Let's go back to the step where we defined the brain region and switch to this other strategy for defining the brain region. So this literally just happened in my lab with a project we're currently doing.

We went all the way to the end of the project. We didn't find the evidence we were expecting. We thought, I think we made a wrong choice back when we decided how to find the part of the brain we were interested in. We went back, we swapped that method. Now we find it in a different way. We go back through all the other steps. And now at the end, we do find evidence for hypothesis.

OK, so now at every one of those choices that I made, all the thousands of choices that lead to the end state, any one of them I could go back to and change my mind if I didn't find the result that I was looking for. But say I do find the result that I was looking for. I'm not that likely to go back to all those other choices and try the opposite version. So if the first way we defined the brain region had given us evidence for our hypothesis, we would have stopped there.

So now the scientific evidence I show you at the end in favor of my hypothesis, we found it by a search process that included seeing the results. We kept changing our choices until we found evidence for our hypothesis. And that, of course, we didn't actually do in that case, just to be clear. You could do that. And in fact, it is common practice and many papers in the field have this form that the final results you see which seem to find evidence for the hypothesis the scientists were looking for have been selected that way.

And if that seems complicated to imagine, I'll just give you the much simpler version, which is called the file drawer problem. I have a hypothesis. OK. I tested it in one experiment. It doesn't work very well. I never published that. I think of a new experiment to test my hypothesis. I run my new experiment. It doesn't work very well. My hypothesis isn't confirmed. I never publish that. I think of a new way to test it.

I do a third experiment. Nope, still don't find evidence for my hypothesis. I never publish that. I think of a new way to test my hypothesis. I run a fourth experiment. It comes out exactly the way I expected. Now I publish it. Look, my hypothesis was right. Here are my evidence. This is true. But you don't know about all those three prior experiments that I did.

So that's another example of many ways that if we publish the specific path that got us to the answer we were looking for. But we can't be honest about all the other paths we took that didn't give us the answer we were looking for. Then what it looks like, and indeed the scientific literature often looks like this, is scientists are always right about our hypotheses. We have our hypothesis. The evidence always confirms it. And so in what sense is that biased? It's biased because you're not seeing all the other ways we tested the same hypothesis that weren't true.

SARAH

Yeah, but devil's advocate question. Isn't it just more efficient to just show--

HANSEN:

REBECCA SAXE: It depends what you think. It depends what I think and what you think. And each given context explains why the last one worked. So one idea is, well, the last one was right. All the other ones were bad. They were wrong. They were not useful tests of the hypothesis. They were broken in some way. Telling you about all my mistakes and everything I did that was broken.

Say I was trying to study how your brain responds to vision, but I forgot to turn the screen on and then it didn't work. Your brain didn't respond, well, because you couldn't see the stimuli. You don't want me to tell you about that. You want me to tell you all the stupid things I did? I forgot to turn the screen on. The headphones weren't working. The scanner didn't run that day. No, that's not scientifically interesting.

So if you're thinking all those failed experiments or all the failed pathways through the analysis plan were broken, then I shouldn't tell you about them. But if you think the opposite extreme is those were all equally good tests of the hypothesis and the hypothesis is false, but the data are noisy and the noise bounces around. So sometimes the noise goes in the direction of the hypothesis and sometimes it goes in the direction against the hypothesis. If I randomly sample ways of testing the hypothesis, mostly I'll get nothing because that's the right answer.

But if I keep trying, I will finally find the specific way that the noise looks like it's in favor of my hypothesis. So you think when people report statistics and they say these results would happen 1 in 20 times, these results would happen 5% of the time if my hypothesis was wrong. But if I tried 30 times until I found the way, then maybe it's not surprising at all that one of those times I seem to find evidence for my hypothesis.

So the extreme end of the spectrum that you said is the successful one is the true one. Everything else was broken. We should only communicate the true one. The opposite extreme end of the spectrum is there is no truth in your hypothesis at all and you just tried 30 or 40 or 50 ways until you find the one where the noise looks like it's evidence for your hypothesis.

Reality in science is usually in between those two. It's some combination of the things you tried that didn't work were broken and the things you tried that didn't work were reasonable ways of testing your hypothesis. And the actual effect is smaller than it looks in the one you selected, but bigger than 0. That's the usual reality.

And so when we go back and replicate prior findings, a standard finding in this literature, if you go back and try to replicate people's previous work, is that the true effect, the one you measure without all the selection, is about 1/3 as big as the effect that was published. So that's been found in many fields and many papers, many different times. So what do I mean by the literature is biased? The selection process that leads to what we publish makes findings look about three times as big as they probably truly are.

SARAH Wow, that's so interesting. And that just speaks to the importance of transparency, because it's really shaping
HANSEN: our reality. And when you think about those scientific findings when they're put into practices like in clinics.

REBECCA SAXE: Exactly.

SARAH Working with real people, that could actually be quite dangerous.
HANSEN:

REBECCA SAXE: Absolutely. And so both the ways in which we've overestimated effect sizes and the ways in which we've overestimated their generality, that they would apply to everybody or every circumstance, both of those are very dangerous when you want to apply science to the real world clinically or in policy or in pedagogy. I mean, we do science because we want it to be applied. But if we're misrepresenting our confidence or the generality of our findings, then it's really dangerous.

And it's dangerous both for the specific application but also for trust in science. If people try applying our ideas and find out they don't work, they lose trust in the whole endeavor. And that's devastating for society. We need shared trust in science and in scientific institutions because scientific process actually often does yield true and useful results when it's working. And those are so important. The true things scientists can discover are so important to health and society and industry that undermining trust through misrepresenting the other things we're doing feels incredibly dangerous.

OK, so you asked. That's the wallowing. We just wallowed.

SARAH I was right there.

HANSEN:

REBECCA SAXE: It's good to wallow. And I want my students to wallow, because if they don't wallow, they could keep doing the current practices that are not OK. But then once you've wallowed, you don't want to stop there, because you just leave depressed and not trusting scientists. And that would be bad. And so then the second half of every class has to be, great, what do we do about it?

SARAH OK, let's go.

HANSEN:

REBECCA SAXE: And so in this particular case, for biased overestimates of our confidence, there's a whole bunch of different strategies, one of which is called pre-registration.

SARAH Oh yes, I was curious about this. I've never heard that term. And you use it quite a bit in on your website and in

HANSEN: your course. So what is it?

REBECCA SAXE: So pre-registration is an effort to be explicit in advance about the path you're going to take through your data so that you specify for every choice you're going to make what choice you will make. That way if you get to the end, you make all those choices the way you planned, and you find the result you hypothesized, then great. That was the test of your hypothesis.

But if you had to deviate, that is, if it didn't work that way you planned and you had to go back and make a change, you now can be exactly transparent about every time you did that, how many times you did it, and with respect to what. And so when you said, can't we just be perfectly transparent about all those times we went back and changed our minds? Well, part of why we couldn't is that we didn't have a practice of being explicit in advance about what our plans were. And it was incredibly hard to know as we were making those changes what changes we had made in response to what evidence and how many of them were there.

SARAH So pre-registration is a new thing? This hasn't been standard practice forever?

HANSEN:

REBECCA SAXE: Well, it's funny. On the one hand, being explicit about your hypothesis before you test them is part of the scientific method you get taught in middle school.

SARAH Yes, I remember that.

HANSEN:

REBECCA SAXE: Oh, I hated middle school.

SARAH I know. Let's have another episode about that.

HANSEN:

REBECCA SAXE: Great. Happy to have an episode about that. But meanwhile, the idea that you could be explicit about your hypothesis and your analysis plans, like for example, everyone who's ever written a grant saying, here's the science I want to do, they had to do that. They had to say, these are the hypotheses I plan to test. These are the methods I will use.

And the oldest version of this that's described as registration or pre-registration in American life sciences research is registry for clinical trials. So that's been a practice for a long time too, of saying if you're going to do a clinical trial of a drug, you have to say in advance, what's your sampling plan? Who will you test on this drug? How many people will you test? What outcome are you measuring? How will you define success? What's your control?

SARAH That makes sense, yeah.

HANSEN:

REBECCA SAXE: Yeah, except that it turns out when you look back at the practice that the published papers deviate from those registrations in really consequential ways a lot. It's so scary. You look back and you see, OK, this drug company said they were going to test the efficacy and safety of this drug with these measures. If you look at their data, if they had done that, they would find the drug is not effective and not safe.

But then the published paper has switched measures. No, we decided instead of looking at efficacy after three months, we would look at it after one week. Instead of measuring safety in this way, we'll measure it that way. Instead of looking at all of the participants, we'll only look at the ones aged whatever, 16 to 19 or whatever. They've picked the subset of the data that now the paper in the scientific literature looks like it found effective treatments with few side effects. But if they had stuck with their original plan, you would have had a very different answer.

So at least now the fact that clinical trials have to be registered means we can go back and look retrospectively and see all of these deviations. Now, that way that I'm describing, it sounds profoundly nefarious. And I will say sometimes I think it is profoundly nefarious. Like in the clinical context, when there's tons of money at stake, it really might be nefarious.

But in cognitive science, in my own training, a lot of it was not nefarious because it's actually incredibly hard to tell when you're doing it. I find this hard to describe. But it feels while you're analyzing data that you're just making reasonable choices. Reasonable, well-informed, rigorous choices. That's what it feels like. And so until I started trying to write it all down in advance, even though I'm an expert and I've been doing this for decades, I had no idea how hard it would be to write it all in advance.

And I will say, I have never yet, I've tried many times, I have never yet pre-registered an fMRI analysis that I didn't have to deviate from. Literally never. Even though I'm as much an expert in fMRI experiments and analyses as anybody. It just turns out it's incredibly hard to foresee all the choices and know in advance what you want to do and make all those choices reasonably.

And the only way to do it, so one strategy is pre-register and then you can be transparent. So now I know in a way I never knew before how often I deviate from my previous plans in response to the data. And I can be transparent about that, and that's a huge help.

SARAH So with the pre-registration, I guess I had this concept that in science there were moments of serendipity or brilliance or oh, I see what we need to do now based on years of training. And so from what I understand, that can still happen. You just have to be transparent about how you deviated from your pre-registration.

REBECCA SAXE: That is what many scientists say is I don't want pre-registration to squash all the creativity and fun out of science. And of course, neither do I. So exactly as you said. There's many ways to handle that. One is to just be transparent. We had this stroke of insight. And the other one is to replicate it afterwards. Halfway through this, we had this insight. We suddenly saw what to do. And then we did the whole thing again and in a calmer frame of mind and having known in advance this time and then we can measure it.

SARAH Great. I feel a little better knowing that the inspiration and the creativity is still in science.

HANSEN:

REBECCA SAXE: Absolutely. No the point is to retain the fun of science and the sense that we can believe in what we're doing. The other thing you can do is replicate yourself. And that always was part of the practice and still is. So realizing, OK, this data set, I've got my hypothesis. I now know what I think happens. I'm going to just test it in new people and be sure. That was part of my practice all along. But pre-registration helps me be more aware of when I need to do that and how much I need to do it.

And part of why that's important is fMRI research is very expensive. And in a lot of scientific fields, people have been resistant to replication, that is to doing the same experiment again, because it seems like a waste of money at the time. We've done one study. It already costs tens of thousands of dollars. Do we really want to do that same study again? Shouldn't we do the next study? That argument that it's too expensive to do the same study again is one reason why many people didn't.

And in every local decision, that can feel right. Really, we want to spend tens of thousands of dollars on just doing the same thing again? But the answer to that usually should be yes, because if you don't, you don't know, you really don't know that the first set of results, how big the effect is, how right you were. You don't know. And the result was so many findings in the literature that weren't true. And that's also really expensive. Publishing and then other people reading and trying to build on claims that aren't true was a huge waste of young scientists' lives, of public money.

So anyway, three solutions to the biased literature. One is replicate experiments. One is pre-register your analysis and be transparent about the choices you made. And the third one is a new publishing model where papers are selected for publication before the results are known. This is called registered report format. It's becoming increasingly popular. The idea is I tell you my hypothesis, my experimental design, and my analysis plan. And if you agree that's interesting and rigorous, you agree to publish it no matter what the results are. And that removes the incentive for me to search through all possible ways of analyzing my data to find the most clean story that will get published.

The thought is by increasing explicit pre-registration so scientists can be more transparent, incentivizing replication so we can tell how true our results are, and then publishing results based on their design and analysis plan rather than their outcomes. All of those shift the practice and the incentives for scientists so that we can be more accurate about the effect sizes and our confidence when we communicate about our results.

SARAH And this just becomes even more important as we're moving into generative AI that's scraping studies and data.

HANSEN: If machine learning is making decisions based on studies that aren't accurate.

REBECCA SAXE: That's a scary thought.

SARAH Yeah, that could be a huge problem going forward. So it makes your class even more important.

HANSEN:

REBECCA SAXE: Thanks.

SARAH And I hope OCW lots of people and lots of educators make use of the resources.

HANSEN:

REBECCA SAXE: Can I tell you about my favorite week?

SARAH Yes, please.

HANSEN:

REBECCA SAXE: My favorite week is the week called, but the incentives.

SARAH Oh yes.

HANSEN:

REBECCA SAXE: So all the way through the class in every previous week, when we talk about data sharing or pre-registration or publication models or science communication, in every week, we would say, here's an ideal way to do science. And somebody would say, but the incentives. And so what did that mean?

The people in the class was, in those years, senior grad students and early stage postdocs would be saying, I want to do science in a way that reaches to a broad audience and that communicates transparently and that allows for open scrutiny. But I also want a career in professional academic science. And the incentives for academic scientists are not aligned. They would say that what I'm going to be evaluated on is how many papers and how high impact a journal I can produce in how short a time.

And so any time I spend on sharing my data or making my code reusable or doing a pre-analysis plan or communicating transparently, any time I spend on that is time I'm taking away from winning this incentives game, which is a competition over a super scarce resource that almost nobody wins and that I can only win by playing by the old game.

And so many people really cared. But there was also a narrative of helplessness, that those of us in the room, grad students and postdocs, we can't do anything about this. The academic incentives, this exogenous thing outside us, the big bad ogres, they control our behavior by incentivizing what we do and what we don't do.

And so I figured I had to have a weak on the incentives. And like every week, it had to have the same form. That is the first half would be wallowing and we would wallow. But in the second half I had to say, there are things you can do. There are tools we can use, there's practices we can engage in. We can change this. That's the whole premise of the class. You have to leave every week feeling empowered. So it was fun to think about. What would that be?

SARAH I'm so curious.

HANSEN:

REBECCA SAXE: How do you empower a class on the topic of the incentives? And so I thought about that. What are the tools that constitute the incentives and how do I empower this room full of people in that context? And so we did a bunch of things. We've done different things in different years. But part of it, as I said, look, the incentives, what is socially valued in science, that's us. That's those of us in the room. What do we socially value? And so we get to decide that and we get to put it into practice.

For example, one of the incentives is who gets hired to prestigious academic departments. My department is a prestigious academic department. So the first year, I assigned my class to read the job ad that was put out for my department that year and propose how they would like to change the language of the job ad so that open science would be welcomed by applicants to that ad. And they made a bunch of proposals and I took those to my department leadership. And the next year, some of that language was in the job ad.

SARAH That's fantastic.

HANSEN:

REBECCA SAXE: Right?

SARAH Yeah.

HANSEN:

REBECCA SAXE: This is stuff we can do.

SARAH Yeah, I love that.

HANSEN:

REBECCA SAXE: Thanks.

SARAH I mean, I feel better.

HANSEN:

REBECCA SAXE: Great.

SARAH I feel like I've been through the wallow-empowerment cycle now. I can see how this course really works.

HANSEN:

REBECCA SAXE: Week after week after week.

SARAH It must be a roller coaster for you.

HANSEN:

REBECCA SAXE: It is.

SARAH Yeah. Well, let's tie it all back to the Saxe Lab. What would you like to do to increase the transparency or

HANSEN: robustness of science? I know you're doing all you can. But aspirationally, what is the next step for your lab to reach this more open level?

REBECCA SAXE: It's interesting because, I mean, obviously there are many practices that my lab can continuously get better about. And we do. We are trying. It's always a trade off. But the main efforts in my life on open science are not in my lab. I mean, as in I encourage the people of my lab to think about data management and resource sharing and broad communication and openness to scrutiny. We try to think about that. But where I see the biggest potential for change or the biggest need for change is at the institutional level is higher than my lab. So in my department and in the School of Science and at MIT more broadly.

So, for example, and again, this partly came out of the class. As I said, in the incentives weeks, I was thinking we get to decide what MIT values. How would we do that? And so partly as a result of that class with Chris Bourg, the director of the libraries, we founded the MIT Prize For Open Data. And so that is a financially remunerative but also reasonably high profile prize for people at MIT who are doing the work of making their science open.

And it's at the whole institute level. It's an incredible event, by the way. People come from across the whole institute and it creates community and makes people see that they are recognized and valued for this thing they were doing because of their personal values. Most of the people who worked really hard to make a data set open or to use an open data set in a creative way, they were doing it because it aligned with their aspirations for science. And maybe they feared they wouldn't get recognized for that, or that there weren't other people who valued them for that.

So to make a big institute prize and have institute leaders there and to give them big checks feels amazing that we can say we do value this. We see you, we appreciate you, and also help connect you to a network of other people who feel the same way. Yeah, so I think that the things I'm doing now are focused on with respect to my aspirations for open science are mostly at the level of the school and the institute rather than my lab.

SARAH Yeah, that's so interesting. I'm curious if there are ways that everyday people like me can support robust science
HANSEN: without being in a lab or without being part of an academic department. How can we support open science through our consumerism or reading?

REBECCA SAXE: That's fun.

SARAH How we engage with our communities.
HANSEN:

REBECCA SAXE: I love that question. It immediately makes me think of a kind of broader value of intellectual humility and caution with oversimplified or over grand or over confident claims. That approach to science that I really value in scientific training and scientific communication feels connected to an approach to anything. I mean, the public funds all of science. So insofar as you feel ownership of the government funding science in your name, I think endorsing and supporting the requirements for openness and transparency, that feels important.

I mean, you are not only a person. You're also a member of MIT. MIT has, as an institution, made a bunch of brave choices around openness, including founding some of the first open registries and going out of contract with Elsevier. There's a bunch of institutional choices that MIT made. And having people who feel affiliated with MIT support and value that, stand up for it, I think helps a lot, helps make the people who made that choice feel empowered to keep doing it.

SARAH Well, I'll tell you what I think we need.
HANSEN:

REBECCA SAXE: Oh yeah. Thanks. Great.

SARAH I think as a citizen, it would be super helpful to have a handy list of questions to, let's say, ask my doctor when
HANSEN: she refers a study to me. I could ask, well, how transparent?

REBECCA SAXE: That's so interesting.

SARAH Was this study? And is it replicable? And how do you know that? Or if I'm at my community meeting and there's a
HANSEN: water study, I could ask, well.

REBECCA SAXE: I love that.

SARAH But I don't necessarily have that in my back pocket, but scientists could conceivably make a list of questions that
HANSEN: as a citizen, I could leverage to be one voice saying open and transparent science is important. And if I'm asking
and other people are asking and all of a sudden all the consumers are kind of demanding that kind of science,
then we could see change.

REBECCA SAXE: I love that. That's fabulous. Yes, please.

SARAH Let's put it on OpenCourseWare soon for everybody.
HANSEN:

That was Rebecca Saxe, instructor, for course 9.401, Tools For Robust Science, which you can find on MIT OpenCourseWare. As always, her coursework is openly licensed so you can reuse and remix her materials in your own teaching and learning. You can help others find the materials too, by subscribing to the podcast and leaving us a rating and review.

Thank you so much for listening. Until next time, signing off from Cambridge, Massachusetts. I'm your host, Sarah Hansen from MIT OpenCourseWare.

MIT Chalk Radio's producers include myself, Brett Paci, and Dave Lishansky. The show notes for this episode were written by Peter Chipman. Jason Player made our YouTube cassette and Shiba Nemat-Nasser built the course on our website. We're funded by MIT Open Learning and supporters like you.