BRAIN SCIENCE PODCAST

With Ginger Campbell, MD

Episode #57

Interview with Neuropsychologist Dr. Chris Frith, Author of *Making Up the Mind: How the Brain Creates Our Mental World*

Aired May 8, 2009

[music]

The *Brain Science Podcast* is part of sciencepodcasters.org, the website where you can find high quality science podcasts from a wide variety of fields.

[music]

INTRODUCTION

This is Episode 57 of the *Brain Science Podcast*, and I'm your host Dr. Ginger Campbell. Today I will be talking with psychologist and author Dr. Chris Frith. Dr. Frith is an emeritus professor at the University College London Wellcome Trust Centre for Neuroimaging. His most recent book is *Making Up the Mind: How the Brain Creates Our Mental World*, which was published in 2007.

Before I start the interview, I am proud to announce that sometime in early April the *Brain Science Podcast* passed one million downloads, and we are currently averaging about 70,000 downloads per month.

Also, I want to thank Audible.com for its ongoing support of the *Brain Science Podcast*. New members can get a free audiobook download at

audiblepodcast.com/brainscience. And don't forget to visit the *Brain Science Podcast* website at brainsciencepodcast.com.

I want to get right into the interview, so be sure to stay tuned after the interview for a few closing announcements.

[music]

INTERVIEW

GC: My guest today is Dr. Chris Frith from University College London. Welcome to the *Brain Science Podcast*.

CF: I'm happy to be here.

GC: Can we start out by just letting you tell my audience a little bit about yourself?

CF: Yes. I am a neuropsychologist, which means I study the relationship between the mind and the brain. My early work was mostly on studying schizophrenia and trying to understand symptoms like hallucinations and delusions. But I was very lucky that in the 1980's I got the chance to start doing neuroimaging right at the beginning of the technological developments. And so, I was in right at the beginning of some of the studies and that enabled me to look at things that were relevant to schizophrenia, like perception, will, and consciousness.

GC: That work is what led you to writing your book, *Making Up the Mind*?

CF: That's right. Yes, that was really an attempt—as much for me as for anybody else—to try and put what I had been doing into a proper framework and understand why it was relevant; to everybody, really, not just to neuroscientists.

GC: And I think you did a really good job of that. In your book, *Making up the Mind*, you started out by talking a little bit about what it means to be a psychologist and how that has changed. I thought maybe you might start out by defining psychology and then explaining how brain imaging really has changed your field.

CF: Yes. Psychology has always been a problematic subject because it falls in the gap between the sciences and the humanities. I was brought up as an experimental psychologist studying it as a scientist. Interestingly, in Cambridge you could either do it as part of natural science, or as part of moral science—which meant philosophy. And there was, of course, this strange school of psychological behaviorism that grew up at the beginning of the 20th century in which you were only allowed to study behavior and you weren't allowed to think about what was happening in the mind—although this has always been a major part of psychology, and we eventually came back to that at the end of the 20th century.

The problem was that this key topic of psychology was actually subjective experience, and this was always considered by other scientists to be not quite right because you couldn't measure it properly. And I guess what I find interesting about the brain imaging is that in a sense this enabled us to start measuring subjective experience. So, ironically, although this is highly expensive equipment for measuring physical things, the environment in the scanner is so restricting that you can't really do very much while you're inside it—particularly in the early days—so in a lot of the early experiments people lay inside scanners and thought about things rather than actually doing things.

And what this showed, for example, was if you imagined lifting your finger, you could see bits of the brain lighting up in exactly the same areas that would light up if you actually lifted your finger. And likewise, if you imagined seeing a face, then a bit of the brain would light up the same as what lights up when you

actually see a face. So, in this sense you had a validation, almost for the first time, of this subjective experience. In the past people would say, 'Why do you believe your subjects when they say they're imagining lifting their fingers?' And now you could say, 'Well, I can see it in the scanner.'

GC: It seems like a lot of experimental psychologists—I've interviewed several—have had the experience that this brain scanning revolution has only validated what they had already found by experiments that people weren't willing to really accept because they said they weren't hard enough science.

CF: Yes, I think that is certainly largely true. And this, of course, is very annoying to psychologists—myself included. I said in my book there's a sort of status hierarchy in the sciences, and psychology is at the bottom of this hierarchy because it's the most soft in the sense that you can't measure things and it's all about subjective stuff. And in a sense the big science of brain imaging got us a littler higher up the hierarchy.

As you said, many of the things that the imaging showed were things that we already knew from our behavioral and introspective type experiments. But this somehow was easier for people to believe. And this isn't just other scientists; it's also the media, and I guess the general public. So, suddenly if you do an experiment with a picture of the brain you can get on the front page of the *Guardian* newspaper, where they would never have been interested in the same experiment purely psychologically reported.

GC: What about the relationship between the mind and the brain? Can you say something about that? I know that's a big question, but I think we need to have a starting point.

CF: Yes, sure. Obviously what we believe is that everything that happens in the mind is in some sense created by brain activity. The problem is that no one

knows how this can happen. And it bridges this philosophical gap between physical processes and mental processes, which of course, Descartes tried to keep completely separate. I think we all still have the feeling that physical processes and mental processes can't really be related. And neuropsychology, of course, is desperately attempting to discover how they are related.

GC: That's I guess one of the themes of your book—to show how the physical processes and the mental processes really aren't as different as we think they are.

CF: That's right. For example, one of the experiments[1] we did was to do with tickling. It's been well known for probably hundreds of years that if you stroke yourself it doesn't feel nearly as interesting as when somebody else stokes you. And tickling—you can't really tickle yourself. We took that into the scanner and basically showed that when you tickle yourself there's not nearly so much activity in the somatosensory cortex, which is responding to touch.

This is in a sense just confirming what we already knew—that you can't tickle yourself—and showing that there's this close relationship between what happens in the brain and what happens in experience. But we wouldn't really expect anything else, because if there was not this close relationship then we would have to become dualists and say the mind is somehow different than the brain.

GC: Traditionally there were the psychologists who tried to study the mind while completely ignoring the brain. Are there still psychologists who are trying to do this?

CF: Yes, certainly. There are a still a few psychologists who say that studying the brain doesn't actually tell us anything about psychology at all; but I think they're getting fewer. And there are certainly psychologists who can do perfectly good

---

[1] Blakemore, SJ, Wolpert DM, and Frith, CD (1990) Central Cancellation of self procduced tickle sensation. *Nature Neuroscience*, 1(7), 635-640

experiments without having brain scanners. And I would tell my students that they should be doing something like 10 behavioral experiments for every 1 brain imaging experiment. I think what changed for me was realizing that when I was studying, doing psychological experiments which might be just behavioral or just introspection, in a sense I was still nevertheless looking at how the brain works.

GC: Right. A major theme of your book is really showing—correct me if I'm wrong—but what I got was that our brain is creating various illusions that have prevented us, in a sense, from developing a science that explains how the brain creates the mind. And you break this down into several parts: the distinction between the mental and the physical; the illusions we have about having direct physical contact with the world; and the illusion that our mental world is isolated. Could you give us an idea of how these illusions have been an obstacle for us?

CF: Yes. I mentioned previously about not wanting to be a dualist who thinks the mind is completely different from the brain. But the way our brain works in a sense makes us tend to be dualists, so it's very difficult for us to think about how the mental and the physical interact. And this is partly because the way the brain works is that it hides from us most of the work that it does.

Something like 90% of brain activity never reaches consciousness at all. And so, we don't know about it through introspection. We can only know about it by doing experiments. And that's why, for example, we think we have direct contact with the world. When I see a tree I am seeing a tree, because I'm not aware of all the work my brain is doing to work out from the very crude signals that are reaching my eyes that this is actually a tree out there.

And that has, I guess, delayed our thoughts about how the brain works, because we just are not aware of all that it's doing. The example I give in the book relates to artificial intelligence. In the 40's when computers came into action, people

thought they would be able to build electronic brains—as the newspapers called it in those days—which would do the sorts of things that humans could do.

And they made a very big mistake, because what they thought at that time was that the easy thing for these electronic brains to do would be to perceive the world, because that's so easy for us, whereas the difficult things for these computers to do would be to play chess, because that's so difficult for us. But it turns out—not that long ago—that the computer has been built that beat the best chess player in the world, but they're still very bad at perceiving things, or reading handwriting, or anything like that.

My friend, Daniel Wolpert, has this nice example that you can make a computer that can play chess but no one has really developed a computer that's particularly good at picking up the chess piece and moving it to the new position on the board. So, we get a very strange idea of what's easy and what's difficult from our introspection.

GC: Can we look a little bit at some of the pieces of this puzzle? Let's just start maybe talking about how we learn about the world around us. You said in your book, "Our brain can know things about the world that our mind doesn't know."

CF: Right. Yes, this is becoming very dualist, of course, just talking about it this way, because our mind is also our brain. But I'm distinguishing between the unconscious bit and the conscious bit. And there are certainly nice experiments showing that, using this technique of masking.

So, you can present a face, for example, very quickly—like for 30 milliseconds—and then send something else immediately afterwards; and you're not aware of seeing a face. Or you can present a face with a very fearful expression immediately followed by a face with a neutral expression, and then you're not aware of having seen a fearful face. But nevertheless your brain—or your

amygdala in your brain—responds to the fearful face by generating fearful responses in you without your being aware of what it is that has caused these. That's just one example.

GC: We only know what our brain tells us, right?

CF: Yes. And somebody had this nice idea that it only tells us on a need to know basis. So, most of the time it doesn't tell us much, because we don't need to know about it.

GC: And now we've seen so many of these optical illusions and things that illustrate this. But what about our bodies—our brain doesn't tell us everything about our bodies either, does it?

CF: No, not at all. That's certainly a point I make in the book. With the visual system you could say we do have a very indirect contact with the world through vision because it depends on light reflected on objects, and they're a long way away, and you can't know about depth directly. But we have direct connection with our body because there are nerve cells sending messages all the way from the little toe up to our brain and everywhere else in the body.

But there's this very elegant demonstration recently discovered called the rubber hand illusion[2], which is where you can't see your own arm but you can see this rubber arm which is oriented in the appropriate way so the hand is facing away from you. If the experimenter strokes the rubber arm, which you can see, and the real arm, which you can't see, simultaneously with two paintbrushes, then after a minute or so you have the strong illusion that the rubber arm is your arm—even to the extent that if he then sticks a needle into the rubber arm you will get a very big emotional response to this.

---

[2] Botvinick, M and Cohen, J (1998) Rubber hands "feel" touch that the eyes see. *Nature*, 391(6669), 756.

All the time the brain is creating in a sense a representation of where your arm is, which usually coincides with your real arm. But you can easily trick it into being somewhere else. I think that's the origin of the phantom limbs that people experience who have one of their limbs amputated. They're experiencing the phantom that we all have, but it's much more obvious because there's no real limb to relate it to.

GC: What really struck me about this was the implications of what you wrote. If our knowledge about the world around us and our knowledge about ourselves is created by the brain, then is the idea that the outside world is fundamentally different an illusion?

CF: No, no. I'm saying there's a real world out there and what our brain does is, given all the sensory information we have and all our knowledge we have from before, it creates its best possible estimate of what it is out there. And this estimate is usually pretty good. But it's never the real world. It's always an approximation of the real world, which we can get better and better with experience. But I'm not showing there's no world out there.

GC: I'm sorry, that's not really what I meant. What I meant was that traditionally there's the world and there's us, and we know ourselves directly but we don't know the world directly. It seems to me like what we really have is no direct knowledge of anything.

CF: That's right. I think that's correct. There's always an estimate in some sense as we are predicting what it is, and we are testing our predictions; and we are improving our predictions all the time.

GC: And so, prediction is a really big part of what goes on.

CF: Yes. Prediction; and also making errors. Errors are very important, because you only know that your predictions are wrong from the errors. And then you can make them better.

GC: What about the fact that even though we're clearly embedded in the world, we seem to have the exact opposite experience—the sense that we ourselves really are separate?

CF: Not only are we separate from the physical world, but we're separate from the mental world of other people and feel ourselves as very much independent agents—although we're also very embedded in the social world as well. The interesting question is why is this a good thing. I mean this is presumably advantageous in some evolutionary sense to have our experience like this.

And one way of looking at it is if you think about vision again, you have a picture of the world on the back of your eye—on the retina—and every time you move your eye this picture completely changes, so that completely different bits of the retina have the various objects that you're looking at on them. And this is happening several times a second. If we were aware of that –

GC: It would drive us crazy.

CF: It would drive us completely crazy. So, the brain has developed this system that stabilizes everything. It says there's a world out there which is completely stable and doesn't move about, and we're moving through it. So, by separating us out in this way it makes experience of the world that much easier to take.

GC: And makes us think that perception just happens.

CF: Yes. That's exactly right.

GC: What about what we already believe or know? Does that affect our perception?

CF: Oh, yes. This is the key thing. Perception is a two-way process. This is why I talk about Reverend Thomas Bayes, who produced this theorem two hundred years or so ago[3]. What he's essentially pointing out is that our perception of the world depends on two things: that is to say, the sensory information that's coming in through our eyes and ears, and our prior expectations and our knowledge of the world. And it's the balance of these two that creates what we experience.

His formula tells you how much do you have to change your model of the world given the new evidence that's coming in. So, if you have very strong expectations, that will affect what you actually perceive. In a sense you can't perceive things that you don't know something about already.

GC: Do you want to talk about Bayes? I don't know whether that's possible in just talking, but it is a very important principle. As physicians, if we order a test for a disease that someone is not very likely to have, then we have a much bigger chance that a result that's positive is really what we call a false positive. That's Bayes, right?

CF: Yes, that's right. That's why Bayes is so interesting. If you put Bayes into Google you get millions of hits. And that's because he's used, as you say, by people doing epidemiology, and pointing out when you interpret a test you have to take into account how likely the person is to have that problem in the first place. And if it's very low, your test is probably very uninformative.

---

[3] Bayes, T (1763). An essay toward solving a problem in the doctrine of chance. *Philosophical Transactions of the Royal Society of London*, 53, 470-418.

There's a big argument in relation to breast cancer. Should you screen all women for breast cancer? But that produces too many false positives. And there's now a suggestion you should only screen people who have a genetic risk, for example; which is taking advantage of Bayes' ideas.

But you also find Bayes being mentioned in statistics. My friend, Karl Friston, who invented the way we analyze brain images, uses a Bayesian procedure for analyzing the data we get out of brain scanners. For example, in the structural scans it takes into account what we know about brain structure and the fact that most brains are very similar to each other, so we're only looking for slight differences.

And also, people who study how the brain works suggest that the brain is a Bayesian system that is concerned with making predictions, and collecting sensory evidence, and then looking at the prediction errors to decide what to do next. And certainly learning about the world these days is very much conceived in terms of a Bayesian process where you predict what's going to happen and then you adjust your learning on the basis of these prediction errors.

GC: Could you give an example, perhaps, of how when we are looking at something that this might apply to what we think we see?

CF: Yes. I guess an interesting example is the ambiguous figure. For example, one of the most famous ones is the Necker cube. This isn't exactly the sort of thing that one should be doing in podcasts, I guess, where you have no visual information.

GC: I think most people are familiar with that, but I'll be sure to put a link for that for anyone that isn't.[4]

[4] http://en.wikipedia.org/wiki/Necker_cube

CF: A Necker cube is a drawing of a cube on a two-dimensional sheet of paper. You basically have two squares and the diagonals are drawn in to give you this picture of a cube. Now, the problem with a two-dimensional drawing of a cube is that it's ambiguous, because you can read it two ways. It could be the left-hand square that's the front of the cube or the right-hand square that's the front of the cube.

Now, an interesting question is if you didn't know, what would you think would happen when you looked at this two-dimensional picture? And what happens, of course, is that first of all the brain interprets it automatically as being a three-dimensional picture, because that's simpler—a cube is a much simpler object than this rather complicated set of lines that you have on the two-dimensional page. So, first of all it's prior expectations that this is three-dimensional when it actually isn't.

And secondly, once it's decided that it's three-dimensional there are two alternative solutions to the problem of what you're actually seeing. And the interesting thing that happens here is that the brain, without your conscious control, switches spontaneously from one to the other. So, you stare at it for a few seconds and then it suddenly switches to the other cube; and then it suddenly switches back again. This is an example of this Bayesian system, if you like, saying because there are two equally good versions of this picture I'm just going to switch between them until I get some information that enables me to decide which one is better.

GC: You have those pictures in the book. And you can't make it switch to the other one. I mean I tried; and you really can't.

CF: Yes. These are very interesting observations with the more basic visual illusions. There is one I have in the book—the Herring illusion—where you have a straight line superimposed on a lot of radiating lines, and you can't help but see

that straight line as curved. To convince yourself that it's not curved you have to sort of get out a ruler and show that it's straight.

And what I find fascinating there is however much knowledge you have that this is a straight line, you can't see it as one. So, there are some times when some kinds of knowledge we have are not sufficient to change our perception, but other kinds of knowledge are sufficient to change our perception.

GC: Is there anything else on that you think we need to talk about before we move on to talking about other people's minds?

CF: I think I could give some examples about people who in a sense have a hallucination—these are perfectly normal people—because they have such a strong expectation of what's going to happen. There's a description from an old book which collected such experiences, of someone who was riding his bicycle in the dark. And there was a storm going on. And his friend was in front.

And he heard a loud crash, and he saw his friend falling off his bicycle and hitting himself. This was actually a hallucination. His friend was actually fine. But he was so worried and his expectation was so great, and it was so dark and thundery that he actually saw this happening. That would be an example where your prior expectations override the sensations that are coming in.

GC: I remember when I was in medical school early on we were doing histology looking at blood cells under the microscope, and you couldn't see the different types of white blood cells at all until you knew what to look for.

CF: A good example of that is, as you know, Brodmann was very important for his studies of the brain because he described all these different regions of cortex on the basis of subtle differences that he saw under the microscope. And subsequently Karl Lashley followed up Brodmann's work about 50 years later and he said he couldn't see these differences under the microscope and it was all

nonsense—which I think many people were convinced by. More recently, using more sophisticated techniques, it turns out that Brodmann was actually right. But that's an example of not being able to see things until you know what you're looking for.

GC: And I guess that applies to a lot of things. Because it seems to me like Lashley's name sticks out in my mind as one of the people who did one of the early experiments that suggested that plasticity existed, but it was pretty much ignored because it didn't fit into the current paradigm.

CF: That's right. And now we know it does.

[music]

I want to take a moment to thank those of you who are supporting the Brain Science Podcast *with your donations. If you are interested in learning how you can help, please visit brainsciencepodcast.com and look for the tab at the top of the page labeled Donations and Subscriptions. If you prefer to send a check, be sure to include an email address with your check so I can let you know it arrived safely. I try to send out postcards, but I always wonder whether they reach you.*

Now let's get back to Dr. Chris Frith.

[music]

GC: One of the things with the whole idea of mind, we have to talk about other people's minds and the processes that allow us to understand other people's minds. We've really learned a lot about that using imaging haven't we?

CF: Yes. Again, a lot of this has been learned from purely behavioral experiments but this has been inspired, I guess, by imaging. And interestingly, as

you were saying, a lot of the early imaging just confirmed what we already knew. But part of this came from classical neuropsychology—that is to say, studying neurological patients with lesions in particular brain regions.

One of the early experiments I was involved with was studying the so-called theory of mind—that is to say, scanning people while they had to think about other people's mental states. And we, somewhat to our surprise, found that a number of brain areas seemed to be involved; including, for example, medial prefrontal cortex and temporal parietal junction. And the neuropsychologists studying patients had never—as far as I can see—studied this kind of task.

So, it first came from the imaging that we said we think these brain regions might be relevant to thinking about other people's mental states. And then as a result people started studying patients with lesions in these particular areas and did indeed find specific problems of the kind that you would expect on the basis of the imaging studies. That was an interesting case where it worked the other way around, almost by chance; because they could have studied these patients previously using such tasks.

The other area which works like that, of course, is the discovery of mirror neurons in monkeys. This was an accidental observation where they had electrodes in the prefrontal cortex of monkeys which fired when the monkey made particular grasping movements, and they found that the neurons also fired when the monkey saw the experimenter making the same grasping movements. Everybody became very excited about that—I suspect slightly overexcited. But the implication was that action/observation and action/execution are very closely related. And then people actually did behavioral experiments, in a sense inspired by the mirror neuron studies.

My friend James Kilner, for example, had people making movements while they watched somebody else making either compatible or incompatible movements,

and showed that if you had to watch someone making an incompatible movement it interfered with your own movements—but only if it was a person making the movements, not if it was a robot making the movements[5]. Quite interesting. So, it was a direct demonstration of this relation between execution and observation. So, again, that was an example of a brain study feeding back into experimental psychology, if you like, and people thinking of new studies that they wouldn't have thought of before.

GC: So, the two fields can inform each other.

CF: Absolutely. Also as a result of the monkey studies people did lots of imaging with humans on mirror systems and they rapidly found that there's not only a mirror system for action but also a mirror system for emotion. So, if you see somebody with a disgusted expression you actually feel disgusted yourself, in terms of activity in the insula[6], for example. Again, that feedback into behavioral studies. There are some people in Sweden who have this very clever set-up where they put electrodes all over people's faces so they can measure muscle activity, and they can actually show imitation of facial expressions in that sort of a way without needing to look at the brain.

GC: In your book you were talking about the brain imaging studies of empathy, and also when we see somebody else that is in pain.

CF: Yes. We feel the pain to some extent.

GC: It's not the physical piece of the pain that we feel, but the emotional part?

---

[5] Kilner, JM, Paulignan, Y, and Blakemore, SJ, (2003) An interference effect of observed biological movement on action. *Current Biology*, 13(6), 522-525.

[6] Wicker, B, Keysers, C, Plaily,J, Royet, JP, Galese, V, and Rizzolatti, G (2003). Both of us disgusted My insula: The common neural basis of seeing and feeling disgust. *Neuron*. 40(3), 655-664.

CF: That's right. Another way of looking at it is that it's well known that you can anticipate pain. So, if somebody knows they're about to have a painful experience then certain bits of the brain will light up—which might have to do with anxiety or anticipation of pain—even though the pain hasn't arrived. And in a sense it's the same bits that light up when we know that somebody else is in pain. Because, again, you don't actually need the pain to worry about it.

GC: Yes, that always puts me in a dilemma as an emergency room physician when I'm getting ready to do something that I know is going to hurt the person. It seems sensible to warn them; but then on the other hand that doesn't necessarily fit what we know about how things work. I'm never really sure exactly what the right thing to do is.

CF: I would guess you should warn them, but only at the very last second.

GC: Right. I try to get things over with quickly. I do believe that makes a difference.

CF: I'm sure that's right.

GC: The processes that go on in the brain that allow us to some extent understand what other people are thinking or what they're intending to do, they are very similar to the processes that allow us to model the world in our own bodies?

CF: Yes, I think so. I mean there are two aspects to that. Certainly the process is the same if you're trying to work out what somebody's intentions are—then the way to do it would be to guess what their intentions are. On that basis you predict what they're going to do next or what they're going to say next. And then you test that prediction by what actually happens, and make your estimate of their intention even better. That's exactly the same mechanism that you would

use in your perception of the physical world. And much of this will run without our conscious awareness.

If I could go back to the mirror story, one of the studies I found particularly interesting, that Sarah-Jayne Blakemore did, was touch. We found that if you see someone being touched—on their face, for example—then the bit of your brain that would be activated if your face was touched lights up, even though it's not being touched. So, in a sense you're sharing their sensory experience by watching them. But what is interesting here is you're not aware of this, that it's happening in your brain.

If you are aware of it, in fact you're a rather unusual person. There is a special form of synesthesia which a couple of people we know have, so that when they see someone being touched they actually say, 'I can feel it on my own face.' The interesting thing to us was that everybody actually has this happening to them, but they're just not aware of it. That's an example of how embedded we are in the social world without being aware of it: we're actually experiencing what's happening to other people all the time, but below awareness.

GC: If you're watching something on TV, like say an object is heading for somebody's head, you have a tendency to duck without even realizing it.

CF: Yes, exactly. Again, this is an example that we're actually very strongly embedded in the social world, and this also alters our behavior. But we nevertheless experience ourselves as independent agents who can do whatever we like; we're not really influenced by what's going on. But in fact we are.

GC: That leads me into what I was going to ask you about next, which is this whole experience of being in control. Some people think that the fact that some of these things are happening outside of our conscious awareness means that our sense of control is an illusion. Do you agree with that?

CF: I am very ambivalent about that. Clearly a very important part of our conscious experience is that we are in control of our actions and that we could have done something else if we'd wanted to—it's sort of so-called counterfactual. It's certainly the case that if you look at brain activity you can see changes in brain activity before you're aware of deciding to act. But I wonder if this is in a sense a sort of priming that happens to us all.

Daniel Wegner has some very nice experiments showing that you can be tricked into thinking that you're acting when you're not, or that you're not acting when you actually are[7]. So, again, the brain is clearly creating our experience of being in control of our actions in the same way that it creates all the other things we've talked about. I am not convinced that demonstrates it's an illusion that we're in control of our actions, but I don't know that I have a very good argument about that. What I am convinced of is it's very important that we experience ourselves as being in control of our actions, because that gives us a sense of responsibility.

GC: Absolutely.

CF: And also we have a very strong experience that other people are in control of their actions. Even children, I think as young as 12 months, distinguish between someone doing something deliberately and someone doing something accidentally. That's already an important part of how we categorize people's actions in the world. And that's key. I talk about at the end of the book social cooperation and being willing to punish people who behave unsocially, but only if we believe they're doing it deliberately.

GC: Right. That seems to be an important distinction that everyone pretty much does agree on.

CF: Yes.

---

[7] Wegner, D (2003). *The Illusion of Conscious Will*, MIT Press.

GC: Can you talk a little bit about something you only touched on in the book, but what you called the problem of privileged access?

CF: Right. It seems that we have privileged access to our bodies that we don't have to the physical world. And there's obviously a difference between the access that I have to my own body and the access that I have to other people's bodies. So, when I perform an action I get privileged access to the sensory consequences of this action from proprioception—that is to say, senses of touch, muscle stretch, and joint angle—which I can never get when I watch other people acting. I can get the same visual access to seeing their actions and seeing my own actions, but I don't get this access to this proprioceptive stuff. So, there's clearly a big difference there.

What is interesting is it's precisely the proprioceptive stuff in my own actions that I'm normally not aware of. The brain clearly uses this. It's very important for making our actions smooth and fast and so on, but it doesn't seem to be very important for our experience of action—unless of course it's gone, and then we're in deep trouble. But that's another story. So, in a sense, although we have privileged access to these sorts of things, it's probably not that important for our experience.

GC: That's kind of ironic.

CF: Yes.

GC: You give an example of something that I think is a good one. You talk about the example of the so-called facilitated communication. Could you talk a little bit about that?

CF: Yes. This is very interesting because it's one of the areas where psychology actually has some sort of consequences in the real world, for once. Facilitated communication is a method that I believe was originally developed for helping

children with physical handicaps who had difficulty moving their hands—who had weakness in their muscles—and it was to help them to use a computer keyboard.

As far as I understand it, a facilitator would stand next to or behind the child and would rest their hands on the child's hands, and would be able to detect the movements that the child was trying to make and would be able to amplify them with their own hands so that the child would be enabled to do what he wanted to do on the keyboard. This then became extended to helping, in particular, children with autism who do not have the motor problems, but often don't have speech, for example. And the idea was maybe they would be able to use a keyboard.

And there were some strange cases where these children who had been completely without speech started writing very interesting things using the keyboard. And then in a few cases more problematic situations arose where they started accusing people of abuse. Some of these cases actually came to court. And the problem was, is the message that you're getting from the keyboard really due to the child or not?

And experiments were done in which the psychologist gave a series of questions which were then supposed to be answered, but they surreptitiously gave different questions to the facilitator than the questions that were given to the child. And the answers were clearly being given by the facilitator. The problem here is that the facilitator sincerely believed that it was the child that was giving the responses. [8]

And Daniel Wegner has done experiments with perfectly ordinary Harvard students showing that you can easily fool people with situations like this into

---

[8] Wegner, DM, Fuller, VA and Sparrow, B. (2003) Clever hands: Uncontrolled intelligence in facilitated communication. *Journal of personal Social Psychology*, 85(1), 5-19.

believing that somebody else is making the response when it's actually them. If you have this very strong opinion that somebody else is going to act and the resulting action is what you expect to happen, then the brain decides, yes, that was them doing it.

Then he has other experiments where he can do it the other way around. You can convince people that they're acting, with things like the Ouija board, when it's in fact somebody else who's making the responses. So, again, if you set up the prior expectation sufficiently strongly the brain will cause you to make wrong decisions about who's actually acting—whether it's you or somebody else.

GC: I guess this has a lot of implications for people that are doing experimental psychology in the design of their experiments.

CF: Yes. That's certainly something you need to take into account. As I say, we're very much embedded in a social context, and we are much more interested in doing what people want us to do than we perhaps quite realize. So, in many psychological experiments there's a conflict between what the experimenter thought the experiment was about and what the subject thinks the experiment is about. And you sometimes may get funny results because of these sorts of misunderstandings. That's another reason why psychology is the most difficult of all sciences, and therefore the most rewarding.

GC: Good point. So, this is what I took away as the bottom line here. Our brains are creating models of the physical world but they're also creating models of the mental world.

CF: Yes. And that includes models of ourselves as well as others.

GC: And as you said in the book, our contact with the mental world is, "…neither more nor less direct than our contact with the physical world. We're using the same processes of using cues from our senses and our prior knowledge."

CF: Yes. Well, there's one big difference between the mental world and the physical world—mostly—is that of course, when I am interacting with someone else, like you for example, I am making models of what I think you are getting at. But at the same time you are making models of what I am getting at. And of course it goes further than that; so I will be making a model of what I think you think I am getting at. So, it becomes quite complicated. Whereas in the physical world a tree out there doesn't give a damn about what I'm thinking about, so it's a one-way process.

GC: Right. And so, the error correction is a lot more complicated in trying to communicate.

CF: Exactly. But the processes are essentially the same. What I find most exciting about the mental world is that once we can interact in this way with other people we can actually improve our models. If your model is better than mine, then I can take advantage of that and update my own in the direction of yours. In a sense that's what science is all about. By getting together with other scientists we've dramatically improved our models of the physical world.

GC: It's important that we realize that's what we're doing—that we're making models—so that we don't become attached to a given model.

CF: That's right. We have to be able to give them up if they don't work. But they are getting better and better. At least I believe that. I'm an optimist. I think our models are getting better and better, in the sense of closer and closer to the actual physical world.

GC: But we don't actually have direct knowledge of anything—not even our own minds.

CF: No, I don't think so.

GC: How does what you have learned influence your definition of the mind? Has it changed?

CF: Yes. I think the thing that's changed most is that I'm more aware of how much what's in my mind depends on other people; whereas no doubt when I was an adolescent I thought it was just me. That's how I've changed a bit.

GC: How do you respond to those who insist that the mind and the brain are exactly the same thing? Do you agree with people who say that the mind and the brain are the same thing?

CF: I think only in part. That is to say I would agree with them that the mind cannot exist without the brain and depends entirely upon brain activity; but along the lines that we were just talking about, because we're so much influenced by the minds of others and by culture. I mean obviously our knowledge of culture comes through the brain, but the existence of culture and of other people makes the mind something much bigger—I don't know whether that's quite the right word—more complex, richer than would be possible just with a single brain. Does that make sense?

GC: Yes.

[music]

GC: If you have a minute I would like to go back and talk about something that I think we skipped over.

CF: Right. OK.

GC: Back on the subject of perception we got off on talking about Bayes and we sort of skipped one thing I wanted to talk about. And that was the importance of both error prediction and acting in the world.

CF: Oh, yes. Perception is not a passive process because, as I said, the only way you can find out about the world is through the errors in your model. And the very best way to do that is to actually act upon the world, so that you say, given my model this is what ought to happen if I do this, and then you can find out whether it does happen and adjust your model accordingly.

It's interesting that engineers have a very different way of looking at the world than psychologists. Psychologists tend to have a loop which says there's perception, signals come in about the world, you interpret them, and then you act. The perception is the input and the act is the output. Engineers look at it completely the other way around. They say you act upon the world, you put something into the world—that's the input, is acting upon the world. And then something happens—which is the output—that enables you to decide what to do next.

And I think this captures this much more active way of thinking about the world which engineers have. Whereas the more passive view of psychologists where somehow you have a perception which you can somehow work out what's going on, it's very much the other way around. We have to act in order to create—to make the world send us back information which helps us to interpret what it is.

GC: And in the imaging work we see, is that one of the ways we know the dopamine systems are involved here? That when we correctly predict what's going to happen, that gives us a good positive dopamine signal?

CF: Oh, no, no. Quite the other way around. The dopamine signal is a prediction error. So, basically if something unexpectedly nice happens, then you get a shot of dopamine; and so, the dopamine neurons become more active. And if you expect something nice to happen and it does happen, there's no response; because there's not an error. If we expect it to happen and it doesn't happen, then the activity goes down. So, that's a negative error.

GC: The thing I find hard to understand about that is, you look at little kids and they like to watch, say, the same movie over and over again where you know what's going to happen. It seems like that would not be a rewarding experience.

CF: Yes, that an interesting problem. I would have to think about it; but you're absolutely right. And in a sense that's true for adults as well—although we don't like to read the same book over and over again, we like to read a very similar book. That seems to be what best sellers capitalize on, or detective stories, or something like that.

GC: We like to be right about predictions.

CF: We like to be right. I have to think about it. It doesn't really fit, I agree. It must be a different system.

GC: Well, there are lots of things left for us to study, then. What's an unanswered question that most fascinates you?

CF: I think we still really don't understand how communication happens—how we are able to share meanings with each other, how we're able to talk to each other. I mean how I'm able to write a book and other people can read it and more or less understand what I'm talking about. This seems to me still a big mystery how this actually works.

And of course a big mystery which I think is not going to be solvable for many hundreds of years is how consciousness emerges from physical activity and the brain. But I think that's so far away that I don't really worry about it. But I think in relation to interactions and communications it will be possible to discover the basic computational processes by which this happens. And I guess the real test of this is in developing artificial systems like robots that can do some of the things that we can do at this level.

GC: Right. Because you have to really understand how it works in order to recreate it.

CF: Yes. But I would also say I am only convinced that I understand how it works if I can create it.

GC: Yes. Is there anything else you'd like to talk about before we close?

CF: No, I think that's fine. You could always come back to me if you want to.

GC: Where can my listeners learn more about your work? Do you have a website or anything?

CF: The functional imaging department where I work has a website which has access to things about my book. But there's not so much additional information there. You can obviously look in PubMed for the various papers that we've written about social cognition and so on.

GC: Yes, because social cognition is one of the areas that you have a lot of interest in. You're an emeritus professor now. Are you still working quite a bit, or spending your time writing?

CF: I'm doing quite a lot of writing. But my wife, Uta—who does a lot of the work with me—and I have a joint position in Denmark part-time with a new project called Interacting Minds. We go there from time to time to take further these sorts of ideas. So, I'm only partly emeritus.

GC: You're still very much involved in what's going on.

CF: Oh, absolutely. Yes.

GC: Well, I really enjoyed reading your book and I hope that this has given my listeners a little introduction to some of the ideas.

CF: Thanks very much.

GC: Thanks again.

CF: Bye bye.

GC: Bye.

[music]

I want to thank Dr. Chris Frith for taking the time to be on the *Brain Science Podcast*. I highly recommend his book, *Making up the Mind*, to everyone. If you are new to exploring the world of neuroscience, this is actually a pretty good place to start. But even if you have read extensively, this book presents a clear description of the current neuroscientific view of the relationship between the brain and the mind. The key theme is that the brain uses the same processes to inform us about the world outside as it does to tell us about our own bodies and thoughts. And in a certain sense all of our knowledge is indirect because it is filtered through the brain.

Yet despite our growing knowledge of how the brain generates the mind, there is an ongoing debate about the relationship between the two. In fact Dr. Frith illustrated the complexity of this relationship because while early on he said our mind is also our brain, when I asked him about this later in the interview he talked about the importance of our embeddedness in the social world and concluded that this makes our minds much larger than the brain.

I am going to explore this idea further next month when I interview philosopher Alva Noë about his book, *Out of our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness*.

Until then, don't forget to check out the other podcasts at sciencepodcasters.org. Vincent Racaniello has an excellent episode of *This Week in Virology* about the swine flu. Dr. Racaniello is a virologist at Columbia University, and he is the perfect person to help us separate the science from the media panic.

Also, don't forget to listen to my other podcast *Books and Ideas* which is available in iTunes and at booksandideas.com. The current episode is the first half of an interview with Dr. Robert Martensen, author of *A Life Worth Living*; and I will be posting the rest of Dr. Martensen's interview later this month. If you are a new listener I hope you will subscribe to both shows.

If you can't support my work with a donation, you can help by telling other people about the show. Links and iTunes reviews are particularly helpful.

Next month's *Brain Science Podcast* will be coming out on the first Friday in June because I will be going on vacation in the middle of June. I will be flying out to Seattle, Washington on June 9th and I will be in Alaska from June 11th through June 19th. I should be back in Alabama on June 20th. I plan to leave my laptop at home, but I hope to be able to get email. If you live in Seattle or Alaska, please drop me an email so I can pick your brain for advice.

As always, you can post your feedback in the Discussion Forum at brainscienceforum.com or send me email at docartemis@gmail.com. The best place to find the links to everything I am doing is to go to gingercampbellmd.com.

Thanks again for listening. I look forward to talking with you again very soon.

[music]

Transcribed by Lori Wolfson
All errors or omissions responsibility of the transcriber