

Angus Menuge: Models of Consciousness (Part II)

<https://mindmatters.ai/podcast/ep134>

Robert J. Marks:

I know I'm conscious, but I'm not sure others are. Talking about consciousness is the topic today on Mind Matters News.

Announcer:

Welcome to Mind Matters News, where artificial and natural intelligence meet head on. Here's your host, Robert J. Marks.

Robert J. Marks:

There is a lot of research happening in modeling consciousness. Panpsychism, quantum consciousness and the integrated information theory are examples of consciousness models that have been getting a lot of press and visibility lately. To talk about consciousness and the models of consciousness, we have returning with us, Dr. Angus Menuge, who is a professor and chair of philosophy at Concordia University. And he's also the past president of the Evangelical Philosophical Society. Angus, welcome.

Angus Menuge:

Thanks for having me back.

Robert J. Marks:

Okay. Before talking about consciousness, it's important to define consciousness. I have been in arguments with people and we go for a long time. And then at the end, we take the time to define the terms we're talking about and find out that, heck we agree, we wasted all our time arguing. So it's important, I think before talking about a topic to define it. So first, what is the definition of consciousness? Is there a widespread agreement to this definition?

Angus Menuge:

Well, the problem is it's an ambiguous term that is used to denote distinct ideas. There is one kind of consciousness, which philosophers of mind have spent a lot of time on, called phenomenal consciousness, which is basically experience, your awareness. So it comes along with the idea of what it is like to see a red rose or to smell that red rose or to feel pain.

Robert J. Marks:

Is this what would be called Qualia. I think you pronounce it different than I do, qualia.

Angus Menuge:

Yeah. Qualia or qualia.

Robert J. Marks:

Qualia, okay.

Angus Menuge:

So the idea, it was once called, they once called raw feels because there is something it is like when somebody steps on your toe or if you get an unexpected check from someone, for example. There's a subjective experience that you have, and it seems to be directly accessible to you. You're aware of it, you can't really deny that you're having the experience. And in some sense, so though some philosophers question this, you have privileged access to it. In other words, we take a dim view when somebody is writhing in pain. If somebody else says, oh, no, you're not really in pain. All right. Because they could be acting, but if they feel that they're in pain, they're not going to listen to anybody else telling them that they're not because they're aware of it directly through introspection.

Angus Menuge:

However, it's not the only notion of consciousness. Ned Block tried to distinguish what he called access consciousness. And here, the idea is more cognitive, it moves from experience to representational content. So for example, if you're solving a problem in logic or mathematics, there is a content to your thinking. That content might not come with any particular qualia or subjective experience. And yet it is accessible to your reasoning. So his idea was that you could perhaps have some qualia that has no particular content. So you just have a vague pain, but it's not a pain that is pointed to anything. And you could also perhaps have thoughts with no associated qualia or experiences or you could have both. So a lot of times when you're thinking about something abstract, you might write something. So you're thinking about prime numbers, but you actually use a symbol to indicate them. So then you would have both at the same time, but they do seem to be distinct.

Angus Menuge:

And then the other kinds of consciousness appear, it seems particularly in human beings, we are also self-conscious so that we are aware of our own awareness. You can, for example, enjoy a sunset, but you can also step back and think about your awareness.

Robert J. Marks:

I've never thought of that being self-conscious is a meta-consciousness, isn't it? That's fascinating. Yeah.

Angus Menuge:

And in fact there seems almost to be no end to the levels of it. This is something actually that Hegel noticed. So for example, assuming that we have good reason to believe that other people have minds, I can, first of all, perhaps I've noticed you and then I'm aware that I'm noticing you. So now I am self-conscious, but then I start to think that you're conscious. So now I'm conscious of your being conscious of me being conscious of your being conscious. And there seems to be almost no end to the levels that you could add. Thankfully, we normally don't. But we in principle can become aware on many, many levels. And maybe one of the most interesting is what the late Lynn Baker called the first person perspective, she noticed that we can be aware, as it were, from the inside of what our life will be like.

Angus Menuge:

So when you're thinking to yourself, "Will I cry at my son's wedding?" That's very different than saying, "Will Angus Menuge cry at his son's wedding?" Or using either a name or a definite description. No, I'm thinking about what it will be like to be me going through that. And that shows, I have an understanding of myself persisting over time. And likewise, when we regret things that we did in the past, or we think about vacations, if such things ever come back again, that we are thinking about what it is going to be like for us to be in those perspectives. And we have a pretty good ability of mental simulation that

allows us to empathize. We can't introspect other people's mental states, but we can to some degree think what it will be like to be that poor person who is suffering now.

Robert J. Marks:

I have this experience all the time. I think so much about my consciousness in this meta state that I don't enjoy life as much as I think I should. I think, "I'm enjoying life," and then I think, "Hey, I'm enjoying life," and I start thinking about my consciousness experience and the entire joy of the experience disappears. It's fascinating. You mentioned qualia. In artificial intelligence I use this as an example of why artificial intelligence will never exist in the general sense where you're going to have a duplication. Qualia for example is our perception of the color red. And I use the example that how are you going to explain the color red to a person that has had no sight since birth? You can't do it. You can explain its properties, its wavelength, that apples are red and other things, but the actual experience is non...it can not be communicated. And if that is the case, how the heck are you going to be able to write a computer program, to explain to a computer what the color red is? Qualia is not algorithmic, it can't be computed.

Angus Menuge:

And that ties in well with the famous example of Mary, going back to Frank Jackson, he imagines a woman, Mary in a room where everything is black and white and she is black and white as well. And she has studied and knows every scientific fact that there is about the physiology of color vision. Trouble is she's never actually seen anything red. And then one day she leaves the room and for the first time sees a red rose. It does seem that she has acquired some new knowledge. She knows now what it is like to see red. And it's interesting. One can get around things indirectly. So colorblind people can stop at stop signs, even though they don't have a red quail because they know what the function of that stop sign is. And they can in a sense, talk about red things and they know what somebody means in a sense when they say that blood is red, for example, but they don't have that same direct, intuitive understanding as the person who has actually seen red.

Robert J. Marks:

One of the evidences of near-death experiences is people who are blind from birth, they have the ability in their near-death experiences to go outside of their body and actually see. So they experience qualia that they have never experienced before in their life. I find that fascinating and really a strong evidence of the mind body problem of dualism.

Angus Menuge:

Yeah. Blind near-death experiences are absolutely extraordinary because they recount information using color terms for colors which they have never actually seen with their eyes. And that's quite extraordinary because it seems as if they had some kind of independent access to them, because it's a difficult question. How could we know what was it like to have that experience? That's an almost unanswerable question I suppose, but it is remarkable that they can recount things using language that describes things which they have never witnessed.

Robert J. Marks:

Okay. Let's get back to some of the models of consciousness here. You mentioned this in the last podcast: Panpsychism. This seems to me to be a cop-out to people that can't define consciousness in materialistic form.

Angus Menuge:

Yeah. Panpsychism does seem to me a rather desperate move. It wants to say that within all of matter, it either has a mind or in panprotopsychism, that it's incipiently mind-like, and that therefore the mind is somehow a potentiality that's built into matter. And it's just a matter then of getting the right configuration and you will get all the wonders of mind appearing. One of the problems with this though, is of course, the unity of consciousness, because if these individual particles are mind-like, and then they formed together, what you would predict and expect is the emergence of many consciousnesses. And in fact, we find the most striking fact about consciousness is that it's unified. So that problem, which is also a problem for physicalism, because I mean, physicalism has this very complex brain. And we now know for certain that the different parts of the brain are used for processing information about different parts of an object.

Angus Menuge:

And yet in consciousness, that object is one thing, like a blue bowl. It's not as if there is a consciousness of blueness and a consciousness of being a bowl and they're separate from one another. There's this objectual unity. And I think that that combinatorial problem is a strong problem for panpsychism just as it is for materialism.

Robert J. Marks:

Yeah. I think the idea of assigning a consciousness to matter the same way you assign mass or energy or something like that is really stretching things. So there are still people that are really backing the concept of panpsychism. And I suppose if you're a materialist, you don't have a lot of options, but that's one of your options. Another model of consciousness is so-called integrated information theory. I had a chat with one of my mathematical heroes, Gregory Chaitin, and we talked about this and I confessed to him I did not understand integrated information theory as being popularized today by Christophe Koch. He wasn't the originator, but he's the popularizer of it. And he admitted to me, and I was surprised, he says, "Yeah, I don't understand it either." This brilliant man didn't understand integrated information theory. In his case, he probably hasn't dug into it as much as he could. Do you know anything about integrated information theory? And do you have any opinions on it?

Angus Menuge:

Yeah, I've looked at it. It's a somewhat interesting approach. It admits the hard problem of consciousness, namely that from nothing we know physically can you predict or explain consciousness. So it suggests that we go about it in the opposite direction. What it basically says is that we first do an analysis of the essence of conscious experiences. And we call them in the theory, the axioms, this is where we're going to begin and we're going to accept consciousness as it presents itself. Now that side of it, I think is admirable. I get disturbed by eliminative materialists like Paul Churchland who seem to deny that we're really conscious, that we even have beliefs and desires, but these phenomena are there and that's denying the facts. So he starts by accepting that there is an accessible intrinsic character of consciousness. And then from that tries to infer, well, what would the physical correlates of consciousness be like to support these characteristics of consciousness?

Angus Menuge:

So it's like a reverse engineering project. And what's interesting too, is that it wants to be a scientific account, it wants to make scientifically testable claims about what the state of the cortex would have to be in order for you to have a conscious experience. And the idea is that it's correlated with the ability to

have integrated representations of a certain kind and so that when you're comatose or drifting off to sleep, what's happening is that ability to form those representations breaks down. And that's the point at which consciousness breaks down. So I think it's worth following and looking into it. I tend to think though, there are going to be some obvious problems with it. It is offering in effect an allegedly causal account of consciousness. But the problem is that there is nothing about those physical substrates that really gives you any reason to expect subjectivity to arise. And there is nothing about those states that really explains intentionality.

Angus Menuge:

So you'll see, sometimes it will talk about the structure or even the geometry of these representations. I'm not sure what's being said anymore, because it seems like there are now physical metaphors being used of our thoughts. So when I think about a triangle, my thought is not triangular and intentionality, it really doesn't reduce to anything physical for some fairly obvious reasons. I can think about the future, but the future cannot be physically causing me to think about it. I can think about the Eiffel Tower right now, and it's not closely influencing me. And I can also think about non-existent objects like elves and hobbits. So the difficulty is even if you could find some of these causal correlates, most likely they are just preconditions. It may very well be that normally if your brain is not in a certain state, you won't be conscious of various things. That's the kind of thing I would expect scientists to be able to give good evidence for.

Angus Menuge:

But there's going to be a gap between these causal preconditions for you to be conscious and explaining what it is that you are thinking about or what it is that you are feeling. There's a content there. And that intentionality doesn't seem to me to reduce to anything physical or be explained by those states of the brain.

Robert J. Marks:

Okay. I think I have a better understanding now of integrated information theory than I did before. I read a report that Christophe Koch gave his theory of integrated information theory to an audience of computer programmers who were very hopeful of a future of artificial general intelligence. And they did not like Christophe Koch's claims that this would be not computable in the near future, that we had a long way to go into development of the future. So that's rubbing people the wrong way, I guess, in some cases. Okay. Another model of consciousness of which I am aware of so-called quantum consciousness. I'm really interested in this because reading the works of Roger Penrose, he maintains that humans can do non algorithmic things. And he looked around at the entire universe and he says, where do things happen in our universe that are not algorithmic?

Robert J. Marks:

And his conclusion was only in quantum mechanics, when you have a collapse of a wave function to a specified outcome, do we have something which is non algorithmic. So I don't know if this relates to quantum consciousness, but there is a theory and a lot of work done in that area. What's going on in quantum consciousness?

Angus Menuge:

Yeah. So the idea of quantum consciousness is that quantum phenomena don't seem to develop in the same deterministic or algorithmic way as things in classical physics. And that this might explain human

creativity and freewill and other powers of the mind which seem to be incompatible with classical deterministic physics. So one view in this area, you mentioned, Penrose, his work is rather speculative because he's looking at quantum gravity and those ideas have not really been sorted out and resolved to this point. But Henry Stapp, following a particular interpretation of quantum mechanics takes the view that perhaps what's going on is that the brain is a quantum system at the level of the ionic activity. And what that means is that there can be a superposition of possible states of the brain. Each one of them, for example, could represent a template for a different action.

Angus Menuge:

So you're deciding, let's say, which of five movies to go watch or watch at home. And there they all, these templates exist in superposition. They all have a certain probability of being selected, but no one of them has been selected. What is it that explains why in the end you watch one movie rather than the others? Well, going back to Von Neumann, Von Neumann had the idea that what's remarkable about quantum physics is that it seems that the observer makes a difference to the evolution of the system. So you can have this system where you have all of these possible states and you've got this wave function. What is it that makes the wave function collapse? Why is it that one of these states actually becomes actual? Well, Von Neumann suggested that maybe it's the act of measurement. Now he himself didn't distinguish between a mental act of measurement or an aura machine doing the measurement, but Stapp does, Stapp speculates maybe the brain is a quantum system and what consciousness adds is selective attention.

Angus Menuge:

So when you're thinking of five things that you can do, the one that you end up focusing on and selecting is fixated. And then that ends up being the one that is realized and you end up actually doing. So, perhaps it is as it were that your mind measures your brain and that your consciousness causes this collapse of the wave function. And that goes on to explain the particular action that you do. And that would be compatible with a very strong view of free will called libertarian free will because no physical state of your brain determined what you were going to do next. It was just your conscious attention that really decided in the end, which of those possible actions that you did, that you weren't simply robotically forced to do it by states going on in your brain as in vagueness system.

Robert J. Marks:

Is quantum consciousness rooted in materialism? Can you look at a materialistic model of consciousness, appeal to quantum consciousness and say, this is materialistic?

Angus Menuge:

Gosh, well, that's a tricky question. Most materialists, their paradigm is really set by older 19th century views of physical science and so by definition, this goes beyond that. However, of course, if one defines materialism in terms of the latest theories of physical science, then you could say that, well, if physical science starts to allow around for consciousness, then I can embrace it. But notice what it does: it will end up in a way trivializing one of the big debates between dualists and materialists. Because if we allow that consciousness is something in itself, Sui generis-

Robert J. Marks:

Sui generis?

Angus Menuge:

Not reduced, with anything else. In other words, it is something of its own type or genus. So it's analogous to, in the history of physics, right, when they thought that electromagnetic radiation required the medium of the ether, and then you had the Michelson–Morley experience that showed that, no, it doesn't require that, it's its own thing. And we no longer regard electromagnetism as somehow reduceable to something that's mechanical. Well, likewise, what if physics will conclude finally, yeah, this is just hopeless. We can't reduce consciousness to any ordinary physical phenomena, but we just recognize it as its own kind of thing. And in fact, we need it in order to have a complete physics, because after all, if you want that theory of everything that Stephen Hawking wants, in the end, as Thomas Nagel said, the theory of everything has to include the scientist as well as the world the scientist observes.

Angus Menuge:

Well, if I am going to have an account that fully explains what's going on when a scientist measures a system in quantum physics and deals with entanglement and all these other things, what if it turns out that that account must appeal to consciousness, does consciousness then become part of physics? If it does then in a way the debate between physicalists and dualist dissipates because the physical has just absorbed consciousness. But the dualists would have won in this sense that they would have cried uncle and admitted that yeah, consciousness doesn't reduce to any of these other things, which is what they'd been claiming for a few centuries.

Robert J. Marks:

Here is the big AI question. This is what I'm interested in. I know that I am conscious. Is there a way we can test for consciousness in others? And if we can, could we apply this test of consciousness in others to artificial intelligence? Can I test for consciousness in you? How would I do that?

Angus Menuge:

Well, it's a difficult question, but it begins I think, with how we are going to generalize on the basis of our data. We find that all individuals naturally, as they develop as children, they develop the theory of mind and that leads them to naturally believe that other people have minds like they do. We are also aware that we do have a mind directly through introspection, and we can see that other people are relevantly like us in every other respect. So it's very reasonable to conclude because it's our natural judgment, but because other people are like us in every other respect to conclude that they have minds. The problem is that when you move to artificial intelligence, artificial intelligence is so different from human beings that now it is not an obvious or reliable extrapolation. So when I test your consciousness by seeing if you produce pain behavior, part of the reason that that is convincing to me is I'm already convinced that you're the kind of being that could have a mind.

Angus Menuge:

With AI, the problem is I'm not already convinced of that. And because the system is so different than us, we run the problem that it might produce all the same behavior. It might simulate all of the behavior you would expect from someone who is conscious. Surely it's easy to program a robot for example, that says, "ow" and withdraws its hand when it touches something that's hot. It can have heat sensors, and it can be programmed to do all that stuff. But that doesn't give me enough reason to think that it's really in pain. And part of the problem is, is because it is so different from me in terms of its makeup. It's different from me in all these other respects and therefore I'm not confident that it's a reliable extrapolation.

Robert J. Marks:

Yeah. That seems to me to be the problem is differentiating between whether or not consciousness is being duplicated or mimicked. And I think that that would be a hard frog hair to cut.

Angus Menuge:

I think so. Yeah. And it's just an odd situation because theoretically it could be that there is something it is like to be this robot or AI system, and yet we would be in a position of being permanently agnostic about it.

Robert J. Marks:

That is really, again, it's an interesting topic and thank you, Angus. I've learned a lot. Hey, we've been talking to Dr. Angus Menuge, who is a professor and chair of philosophy at Concordia University. We're going to have him back for one more subsequent podcast. And until then, be of good cheer.

Announcer:

This has been Mind Matters News with your host, Robert J. Marks. Explore more at mindmatters.ai. That's mindmatters.ai. Mind Matters News is directed and edited by Austin Egbert. The opinions expressed on this program are solely those of the speakers. Mind Matters News is produced and copyrighted by the Walter Bradley Center for Natural and Artificial Intelligence at Discovery Institute.