

Consciousness, Materialism, and the Limits of Reductionism

<https://mindmatters.ai/podcast/ep316>

Austin Egbert:

Greetings and welcome to Mind Matters News. This week we're wrapping up our discussion with Dr. Selmer Bringsjord about integrated information theory and an alternative view, cognitive consciousness. This is the third part of this series, and I'd encourage you to listen to the prior parts if you haven't already. Now, here's your host, Pat Flynn.

Pat Flynn:

This is interesting, and I want to get this point out. I understand, I guess, maybe some of the wider philosophical assumptions or the wider philosophical backdrop at play here. I mean, if you talk to good old Aristotle or an Aristotelian, the reason they would say that an animal is conscious or we are conscious is because it's tied to certain powers like sensation and locomotion. And then once you go down the chain of being, and those things aren't there, there's really no reason to think that they have consciousness. There's no, on a sort of basic point, but then if you come, I guess from a general materialist standpoint, and Aristotle is not a materialist, so I don't think that he has this problem, I guess the issue here is how, through a series of quantitative steps, you can leap over what seems an enormous qualitative abyss. That seems to be the fundamental issue that is constantly recurring. Right?

Selmer Bringsjord:

Well said.

Pat Flynn:

Because if you are a materialist, you think that what is at bottom is essentially the exact opposite of everything that our conscious life is. It's not intentional, it's not directed, it's not about anything. It's not teleological. There's no, what-it-is-likeness, right? None of that is there, right? So how do you take whatever that stuff is that is totally devoid of any of those features and put them in various combinations or permutations, and then suddenly it poofs out, to put it somewhat polemically, this reality that is qualitatively utterly unlike where it just apparently emerged from. I think that's the fundamental issue.

And so I think a lot of naturalists have realized, okay, maybe the problem is that we can get over this problem by not having the gap there in the first place, and we just have to have either consciousness itself or some type of proto-consciousness or something like consciousness down at the ground floor somehow, and then we don't have this sort of gap problem.

And that's an interesting move. I think it's a concession that materialism and physicalism are false, and I think what we're talking about now is no longer anything like that worldview, and it's debatable of even arguing whether it's even a naturalism or not. So talk to us about all of that. I just threw a lot of stuff out there, but it seems difficult to have these types of conversations without trying to make these background wider, philosophical assumptions explicit, and when or when not it's legitimate to bring in something like panpsychism without having to admit that the prior world picture you had is essentially refuted by that ad hoc move, if that makes sense. Yeah, talk to us about that.

Selmer Bringsjord:

Yeah. Yeah. It's an ad hoc move considered from the standpoint of sort of pure rationality, but it's not ad hoc from the rhetorical point of view or the subjective point of view. Pure objectivity, which frankly you're using to minimally cast some doubt on it, if not outright render it overwhelmingly suspicious, says it doesn't make any sense. It rhetorically makes sense. I don't think anyone, so Naveen and I, my co-author in this paper, I think in many ways we both start with an honest observation which Christians and other historical monotheistic religions could get in here as well, but I'm not, I'm a Christian.

So here in listening, for example, now just about done with Marilynne Robinson's book, *Genesis*, which is just absolutely unbelievable, the observation is a two-part one. The observation in terms of simply the behavior and accomplishments of creatures on our planet look. We take the simplest working system engineering-wise that we sort of are semi-impressed with. So most people don't know how a toilet operates. This has been empirically demonstrated. A doggone toilet is actually really a smart system for how the siphon action works. We came up with that, okay? I don't need to talk about Ferraris, but I can. I don't need to talk about rocket ships, but I can. I don't need to talk about space stations, but I can. Using a hammer or using a rock as a hammer and tool use, that's not what I'm talking about. I'm talking about the engineering of a system.

We're just simply, let's be honest, brilliant creatures. We're still doing things like making deep fakes and invading countries for no ethically sound reason, but intellectually we're unique, and we don't know if there, we haven't been to Alpha Centauri, so I don't know, maybe there are aliens that are super smart beyond us. We keep saying, the AI people keep saying, "Yeah, well, we'll make the super smart things by ramping up AI." Yeah, I've been hearing since I was an undergraduate. It's not going to happen anytime soon by any metric.

So we start with the observation, look, we got this big problem here. We got this fundamental divide. What is going on here? Scientifically, formally, mathematically, engineering-wise, what's going on? And I think somehow, just as you say, somehow, a lot of these folks on the IIT side and the PHI side say, "Well, all these other things also are phenomenally conscious, even things that you might say are absurd to consider so like seemingly inanimate objects. We just have a little more integration and so forth."

That's ridiculous. It's ridiculous. We have this thing called math and engineering that Leibniz and Newton were pretty, it's a big deal that we have Calculus.

Pat Flynn:

It is ridiculous, and it parallels, I think, other kind of extreme positions in philosophy. I mean, think of mereology, the study of parts and wholes. You have two extreme positions. You have universalism, which is the idea that any random things you point at form a legitimate object. My foot, Donald Trump, and my minivan form a legitimate object. And I think most people realize, no, that's insane. That isn't a legitimate object. It isn't tightly unified enough. And then there's nihilism, which is the idea that there are no composites at all, everything is sort of erratically simple. And most people realize, no, there, there's got to be some sort of reasonable middle view there, and Aristotle offers it. It's a substance based ontology where we can differentiate between real, substantial unities, accidental unities, aggregates, and so on and so forth. We don't have to get into that now.

But this debate about consciousness, it seems to be that people are, if you're sort of a materialist, it's just swinging from one extreme to the other in a parallel way. There's people who are materialists that are eliminativists, and they realize the same problem. They realize that if we don't have any qualitative dimension at the bottom, we're not going to get it at the top, so it can't be there, and whatever we think this consciousness thing is, it's not what it seems like, it's an illusion or whatever. You know what eliminativists say. And most people are like, "No, that's insane. This is just definitely not correct." And

you can argue against it and give all sort of reductios, but I think some things can really just kind of be dismissed with a hand wave. I'm just being honest about it, right?

On the other hand, it seems like people, the same types of people with a similar paradigm, realized a problem. And so they're just going to the other extreme instead of saying that nothing is conscious, everything is conscious, right? And that's definitely wrong too. So it just seems like a very parallel situation where what we should do is get back to that same middle position and find a world picture or worldview in ontology that allows us to make sense of that. And I would say, and I think you'd probably agree, that something like an Aristotelian view broadly can really help us do that, so we don't have to land in either of these, honestly, crazy positions, and we can still do just as well with all the science. It's empirically equivalent and all this or that. So yeah, your thoughts on all that, it just struck me that, yeah, this just seems like another just similar case of sort of materialists, because of their paradigm, being forced to swing between two various extremes, where there's a very reasonable position that could be maintained if you're just willing to just forgo that materialist physicalist starting point.

Selmer Bringsjord:

I couldn't agree more, no question about it. The eliminativist case has of course been instantiated by, well, maybe sometimes he got sheepish about it, but Daniel Dennett. I think that was, it's all an illusion and so forth and so on. Yeah, I mean, you can call the thing an illusion, but then when you finished writing your book, and you go get a cup of coffee-

Pat Flynn:

You talk about how good it felt to write that book.

Selmer Bringsjord:

Well, there you go. And then if you go get, maybe you celebrate after finishing it with a glass of champagne. I don't know if he engaged in such things like, "Oh, this is delightful, this champagne. So happy to be celebrating my achievement." Interesting self-refuting move that you would make in registering your appreciation of the champagne, so I think that's the other problem with the view, the offshoots of the move that you are pointing to, the eliminativist move. So yeah, no, it's amazing. It's amazing that we have these groups. I respect them. I like to talk to them. Some of them are my friends, but really at the end of the day, goodness gracious, I mean.

Pat Flynn:

Yeah, well, Aristotle himself, a little error in the beginning can be multiplied a thousandfold by the end. I would say probably it's more than a little error at the beginning, but it is interesting to see how these rather absurd implications, and it is funny to see that these sort of panpsychist moves, which of course comes in different models and flavors, are becoming increasingly popular. And I think that that itself is an admission, maybe not an explicit one, of the failure of the materialist paradigm to account for these, for not just the hard problem of consciousness. Consciousness is just one of the issues. I mean, James Ross isn't even focused on consciousness. He's focused on rationality and so was Aristotle. They didn't really think about what it is, likeness. They were thinking about these rational powers that we have, and these are logically distinct issues, aren't they?

Selmer Bringsjord:

They are. And we haven't talked about what might, I don't think you can have multiple Achilles heels. I guess you could, if it's the humanoid case, you can have two. But we haven't talked about another, if

you will, Achilles heel here in the affliction of this camp, intellectually speaking, which is the utter absence of a precise definition of what's being talked about.

So I think it's fine in a conversation like ours to refer to phenomenal consciousness as that which it's like to, as an example of it, experience a fine single growth pour over coffee from your high-end coffee house or the ski turn or whatever it is. That's fine, okay, but that's not how science works. So these people are doing science, and they're allowed, for reasons that I have yet to fathom, including in the recent paper on entropy, they're allowed to refer to consciousness over and over again and pursue it, seek it, and they haven't defined it.

I mean, the lack of definability of the very thing that's at the core of human existence should be a signal to you that something unbelievably profound is going on here. Don't fake it. Don't fake it and give people the impression that it's not that big a deal, we can handle it, we can formalize it, we can mathematize it, and we can look for it empirically. No, you can't. You can't even define it.

And the lack of humility in the face of that, given that, again, this has criticisms made of Dennett, I think repeatedly, let me get this straight. You're telling me the very phenomenon that is the reason I'm staying alive and not just ending my life or that it's the very reason why no one should monkey with my life, *ceteris paribus*, is an illusion? Okay, really?

Well, what is this thing that's an illusion? Can you please define it for me? He never gave a definition of phenomenal consciousness, so you're not allowed to do that. You might be allowed to do that 200, 300 years before Christ, when Aristotle, bless his heart, has the first logics pretty darn good mathematics. This is 21st century. We've got a lot of math, a lot of logic, a lot of theory of computation. If you really know what you're talking about, just define it, write it down, write the definition down, and then we'll go investigate it or step back and say, "Oh my goodness. This is the very essence of why life is meaningful, that I have a mind that can have genuine experience, that I'm conscious, that I have feelings, that I have hopes, and I have dreams, that I can have beliefs." You mentioned intentionality. "That I can direct my mind at remote objects and abstract."

If you can't define it, don't fake it. Step back and say, "Wow, how in the world am I going to explain this?" And good luck to you trying to explain it from a materialist point of view.

Pat Flynn:

Yeah, that's a great point, and we've talked about moral implications, and obviously there's a lot of practical implications here. I mean, ideas have consequences, and these perspectives are often just so seriously diminished and impoverished with respect to morality and meaning. And I know not all atheists or naturalists are nihilists or want to be nihilist, but it's hard to avoid nihilistic consequences, often outright depressing nihilistic consequences, from a lot of these perspectives, especially eliminativist ones.

And I think if you are an eliminativist and if you're trying to be consistent, I think you should be a nihilist in many respects. I mean, think of somebody like Alex Rosenberg. He's got that book, *An Atheist Guide to Reality*. He's an eliminativist, and he's a nihilist. I think he's deeply wrong, but I think he's at least pretty consistent. And I know a lot of atheists, naturalists, they don't like his project, but it's hard to see why, if you aren't committed to a certain world picture, a certain ontology, a certain epistemology, a sort of broad scientism, why you shouldn't be on his program. It's hard to avoid.

And that's sort of like his mission. For anybody who's read his book, *An Atheist Guide to Reality*, he doesn't really argue against God in that book, he just sort of lazily gestures towards Hume. He's like, "Oh yeah, Hume took care of that problem." This is really a guide for atheists. He's like, "If you're an atheist, this is what you should believe." So it's really a book written for atheists trying to get their beliefs

straight. I think it's pretty convincing in that respect, and I think he's sort of an unintentional ally to people like us in the sense that he's giving us aid because the conclusions are so absurd that if they really are tied to that starting point, it should give people really strong reason to go back and revise that starting position on all sorts of fronts, not just with respect to morality and meaning, the things that you've been talking about, but of course to just other things that are just flatly undeniable like the fact that we do have phenomenal states.

There really is what-it-is likeness to have this conversation with Dr. Bringsjord while I'm sipping my coffee, and that the life comes to be through this, that is my world. And to deny that just fundamentally undermines everything, including any sort of practical, reliable, scientific inquiry into the world. I mean, it is so core and so fundamental that I don't think you can coherently, at the end of the day, actually sort of eliminate the sorts of things that this paradigm is trying to eliminate. So this is just kind of my long-winded of saying, I definitely agree with you. And I actually think pointing towards certain people like Bennett and Rosenberg can actually, in a sense, be helpful because I think that they are, I don't think they're stupid, I think they're actually pretty good thinkers that are just trying to work out the consequences of a particular paradigm, and they do it in many ways more consistently than other people.

Selmer Bringsjord:

Yeah, no, I couldn't agree more. Absolutely.

Pat Flynn:

Yes. Dr. Bringsjord, before we wrap up here, is there anything else that we haven't said about integrated information theory that you think is important that other people should be aware of? Any other aspects of your paper or argument? Obviously we're going to include a link. This is available online for free as a bonus to the volume of *Minding the Brain*, so people can read it at their leisure, and we'll encourage them to do so. But yeah, any other final thoughts on anything?

Selmer Bringsjord:

Well, perhaps just one. We haven't talked much anyway, I've sprinkled in a few references and some more specific information about the essence of the competitor here, cognitive consciousness, and that's pretty easy. I could actually just use Dennett, an example of his to perhaps crystallize the theory and the difference. I think Dennett says at some point, "Well, I like dogs." Obviously I'm paraphrasing. "I like dogs. Dogs are pretty smart. Darwin thought dogs were really amazingly smart, even stacked against the human mind. But look, I put the food in the bowl, and the dog believes there's food in the bowl. And the dog saw me do that, so the dog believes that I not only put the food in the bowl, but believes that I believe there's food in the bowl."

So that's two levels. Okay, does the dog believe that the person behind me who's not the so-called master, but always around, another family member, believes that the master put the food in the bowl? If so, the dogs got a third order belief. And Dennett said, "I could go along with that maybe, probably." But we, I mentioned detective stories earlier, we routinely go way beyond that. And cognitive consciousness as a theory says, look, that's the mark of serious levels of cognitive consciousness. And if you get, we believe dogs are cognitively conscious, but now when you work your way down, if we don't have a single layer of cognitive verbs that are operative in the agent, we don't have that the agent knows P for some meaningful proposition. Sorry, there's no cognitive consciousness.

When you apply our measurement scheme, λ , to a pure deep neural network, a pure chatbot, it says the level of consciousness, this would be my demonstration in the cognitive consciousness case for

Blake Lemoine, I'd say, "Yeah, well, you don't know what phenomenal consciousness is anyway, but we could talk more about that."

In the cognitive consciousness case, we can prove that large language model, if it's a pure deep neural network, has zero cognitive consciousness. It literally has no beliefs. You can't find the structures in the system that correspond to an agent believing a proposition. You can only find numerical data swimming dynamically through time. And what John McCarthy said originally, we have an inspiration from John McCarthy, what he said was he was echoing what he said around 56. He was echoing that in 2006 at the celebration of AI at 50 years of age. He was saying, "You folks are talking all about the statistical based processing of these systems, but do you understand they don't actually know anything? And worse, they don't know they know anything. So if they don't know, they know something or know they don't know something. And beyond that, how are they intelligent? They're not." And they're certainly not cognitively conscious. So I would just throw that in there. And of course, many more details along that line are provided in the chapter in question.

Pat Flynn:

Yeah. Well, that's great. That's super helpful. And the paper is, it's great. It's a technical read, so people should be ready for a challenge. But we like that. We like rigor, we like technicality, we like challenge, and Dr. Bringsjord, let's finish with this. What are you currently working on next? What can people look forward to, and how can people keep up with your work?

Selmer Bringsjord:

I don't really know how they can keep up with my work systematically.

Pat Flynn:

Not that regularly on social media like the kids are these days, yeah.

Selmer Bringsjord:

Yeah, they can Google stuff, and they can go to my CV, which is a little behind. But in terms of stuff I'm currently working on, one of those items I alluded to in passing is extending our concept of consciousness and our measurement scheme for it to the quantum computing case. It's going to take a while before we have some publishable results, but maybe a year, year and a half. And that's really exciting.

Pat Flynn:

That is exciting.

Selmer Bringsjord:

I got caught off, I mean, this quantum computing thing, it's one thing to talk, Deutsch long ago, introduced it in concept, the concept of a quantum computer, but by golly, back to the point we discussed about humans being singular in the known universe and making systems, systems of complexity, that are amazing. If a quantum computer doesn't amaze people, an actual quantum computer, and we have one at my university RPI thanks to, well, thanks perhaps to a billionaire in IBM. It's an amazing thing. So I'm really excited along with my colleagues on this, about the exploration of AI in general, intelligence systems in general, and also specifically cognitive consciousness in a quantum computer. I find this entrancing. So hopefully some results coming down the pike.

Pat Flynn:

Yeah. Well, great. Well, we'll keep up with all of that, and hopefully that'll give us an opportunity to have another conversation in the future. I would certainly look forward to that. For all the listeners out there, I would like to encourage everybody to grab a copy of *Minding the Brain* if you have not already. Dr. Bringsjord has really two contributions. He has the development of the Rossian argument in there which is really just absolutely I think it's just such a good argument. A little bit challenging again, but once you grapple with it, I just think it's really powerful. And then this other article, which is free online, I'm sure the links will all be in the relevant places, so thank you all for tuning in. Please do like, share, and review the podcast. We'll talk to you all soon. Goodbye.

Announcer:

This has been Mind Matters News. 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.