

ANNALS OF THE NEW YORK ACADEMY OF SCIENCES

Issue: *The Emerging Science of Consciousness: Mind, Brain, and the Human Experience***The thinking ape: the enigma of human consciousness**Steve Paulson,¹ David Chalmers,² Daniel Kahneman,³ Laurie Santos,⁴ and Nicholas Schiff⁵¹Wisconsin Public Radio, Madison, Wisconsin. ²Australian National University, Canberra, Australia; and New York University, New York, New York. ³Princeton University, Princeton, New Jersey. ⁴Yale University, New Haven, Connecticut. ⁵Weill Cornell Medical College, New York, New York

What is the origin and nature of consciousness? If consciousness is common to humans and animals alike, what are the defining traits of human consciousness? Moderated by Steve Paulson, executive producer and host of *To the Best of Our Knowledge*, Nobel laureate psychologist Daniel Kahneman, philosopher David Chalmers, expert in primate cognition Laurie Santos, and physician-scientist Nicholas Schiff discuss what it means to be conscious and examine the human capacities displayed in cognitive, aesthetic, and ethical behaviors, with a focus on the place and function of the mind within nature. The following is an edited transcript of the discussion that occurred October 10, 2012, 7:00–8:15 PM, at the New York Academy of Sciences in New York City.

Keywords: consciousness; philosophy; ethics; behavior; cognition; mind

Steve Paulson: Welcome. It is wonderful to see such a terrific turnout here. I'd like to say a huge thank you to the Nour Foundation and the New York Academy of Sciences for making this event possible.

It is a great pleasure to be here because we have a terrific panel and some fascinating ideas to dig into. I have been neck-deep in questions about consciousness for the last 4 to 5 months. I'm in the process of putting the finishing touches on a 6-hour radio series on the science of consciousness, which will be coming to a public radio station near you in the coming months.

I'm not a scientist or a philosopher; I'm a public radio guy. But for whatever reason I can't stop thinking about the subject of consciousness, which, in one sense, is rather odd—my wife will be reading a great novel while I am plowing through a philosophical tome about the mind/brain problem; I can't really explain it but reading books about the nature of consciousness is strangely addictive to me.

Let me give you two recent examples. The neuroscientist Christof Koch, who did groundbreaking work with Francis Crick, recently came out with a very interesting book entitled *Consciousness: Confessions of a Romantic Reductionist*. I interviewed Koch and mentioned that some scholars, including the distinguished philosopher on our panel David Chalmers, have suggested that science will never understand certain dimensions of consciousness. Koch replied, and I quote, "If you look at the historical record of philosophers, it's pretty disastrous. Science has a spectacular record of understanding the universe," and he went on to say, "I'm profoundly skeptical when philosophers tell us once again what we'll never know." [*Audience laughter*] Something to talk about this evening, I think.

Take another example, the philosopher Thomas Nagel has a new book called *Mind and Cosmos*, also quite interesting. Nagel wrote the famous essay, "What Is It Like to Be a Bat?" The answer, by the way, is we will never know. Nagel's new book is a critique of the standard materialist model of science and specifically the way many neuroscientists try to explain consciousness through neural correlates. He ends the book by saying, and I quote, "I would be willing to bet that the present right-thinking consensus will come to seem laughable in a generation or two."

So here we have fundamental questions about science and philosophy, not to mention a certain degree of testiness when it comes to trying to explain the nature of consciousness. Of course there are all sorts of other big questions as well; for instance, what kind of consciousness do animals have? Will computers become conscious someday? And what about the people who have fallen into comas after

suffering severe brain trauma; where does consciousness begin and end with them? This is fascinating stuff. We'll be talking about all of this and more on our panel, "The Thinking Ape: The Enigma of Human Consciousness."

Let me introduce our very distinguished panel of speakers. David Chalmers is a philosopher of mind and consciousness at New York University and director of the Center for Consciousness at Australia National University; his many books include *The Conscious Mind: In Search of a Fundamental Theory*. Laurie Santos is a professor of psychology at Yale University investigating the evolution of the mind, the theory of mind, and the development of cognition in humans and nonhuman primates. Daniel Kahneman is a Nobel laureate and professor emeritus of psychology at Princeton University who pioneered behavioral economic theory; he's the author of *Thinking, Fast and Slow*. And Nicholas Schiff is a physician and scientist at Weill Cornell Medical College, where he focuses on the pathophysiology of impaired consciousness, arousal regulation, and the effects of deep brain stimulation techniques on minimally conscious patients. It is great to have all of you here.

Dave Chalmers, let me start with you, since I mentioned you earlier. Some people say understanding consciousness is the biggest mystery left in science. What do you think?

Chalmers: I've always seen it as pretty well the biggest *challenge* for science, for a scientific world view. I started out in the sciences and mathematics and physics, and there are a lot of puzzles in these areas. And from working in the middle of them I got the sense that scientists basically have a correct world view, and that they are currently cleaning up some of the puzzles; that while we're not quite yet at the last stages of explanation, we've got a sense of what the relevant picture of the universe, of what the domain looks like. As a result there's a beautiful scientific picture—a great chain of explanation: physics explains chemistry, chemistry explains biology, biology explains at least some aspects of psychology, psychology explains aspects of sociology, and so on. And although there are a number of details yet to be worked out, we've at least somehow got a sense of the whole picture and how pieces fit together.

What's interesting about consciousness is that it just doesn't seem to fit easily into that picture at all because our scientific picture of the world is described in terms of objective mechanisms from the objective point of view. In contrast, consciousness is the quintessentially *subjective* phenomenon; it's how things feel from the inside; it's how we experience the world from a subjective point of view. But *nothing* in the scientific objective picture of the world seems, on the face of it, to tell us why there's going to be subjectivity.

So I see it—by the way, I didn't say what you said I said to Koch . . . [audience laughter] . . . I never said science *can't* explain consciousness . . .

Paulson: . . . you have hinted at that very strongly . . .

Chalmers: . . . certain kinds of *standard* scientific explanation wholly in terms of brain mechanisms may fail, yes. But I see it more thoroughly as a challenge to science. It may be that our methods of science and our theories of science have to be expanded to bring consciousness in.

For years, I've organized a conference called "Toward a Science of Consciousness," so I'm proscience; I'm a glass half-full guy . . .

Paulson: [laughing] I never said you were not proscience . . .

Well, we will come back to that, to how far science can go in explaining consciousness.

Let me just throw this open to the rest of the panel. Is consciousness one of the big questions out there? Is it one of the big mysteries? Or have we overblown this? Is it not as big a challenge as we're saying? Niko Schiff, let me turn to you.

Schiff: I totally agree. I would say that the science of consciousness is extremely challenging and that in the context of trying to make operational evaluations of patients, that is, when trying to determine if they are conscious or not, we don't have a standard model—I don't even think there's a dogma (the idea of a standard model here is laughable). So, while we do have measurements and some operational approaches,

and while we certainly know in a casual sense when somebody's conscious, it is very difficult to demonstrate that a comatose patient who starts to recover but inconsistently responds is actually conscious. If the patient never responds, could he/she be—or become—conscious? Is the patient conscious now and we just don't realize it? We're getting better tools that provide measures for approximating probabilities of a state of consciousness.

Actually, I would say that as I've looked at this, and my colleagues have looked at this, more carefully with better measurements over the last 10 years, measuring and predicting consciousness are more challenging and harder than we originally thought. I realize how many mistakes I've made along the way, and I still make them. So I find the problem of consciousness very challenging, a very humbling kind of problem to attempt to solve.

Paulson: Laurie Santos, let me turn to you. I know your specialty is animal cognition; so is this a big question for you as well: What is consciousness?

Santos: I agree with what people have said so far. In fact, I don't think we modern neuroscientists and cognitive scientists know how to get at the phenomenon of subjective experience—when does it occur; what does it feel like to have subjective experience; how can we measure it? That said, cognitive science has made tremendous inroads into other things that were once thought unmeasurable. For example, if this discussion were held in the 1950s, a group of behaviorist scientists would be sitting on the panel and they would say that while behavior can be measured fantastically well, the “black box” of the mind is just going to remain a black box.

And yet, since the 1950s neuroscientists and cognitive scientists have come up with all kinds of cool techniques to probe what's going on in the black box, both behavioral and neuroscientific tools.

And so—[*speaking to Chalmers*] I didn't think you'd be the optimist up here about measuring consciousness—I also take a glass half-full approach: I think that while we definitely don't know where to search for an explanation of subjective experience, this doesn't mean that 50 years from now we won't all be in this room saying, “Oh, we have this fantastic tool and we know what we're doing.”

Paulson: Danny Kahneman, how big a question is this in science, consciousness?

Kahneman: Well, that's very odd. I'm in a minority because for some reason I'm one of those people who never got myself completely fascinated by this question [*audience laughter*]. And in part this is because I never could imagine what an answer to that question would be. I find it difficult to conceive of a question without having some idea of the structure of an acceptable answer. If there is a structure, I don't know about it.

What I *do* see—[*to Schiff*] and that's the approach that you talked about—is that in fact we *can* identify consciousness; we can agree on it. And while it is subjective, we can evaluate the consciousness of other people and of other animals, and we're getting better and more consistent at it.

And so building from the bottom up, I think we can get a better understanding—or at least a better description—of the conditions for consciousness. To attempt to bridge the gap between the material and the subjective, I don't know how that gap could be bridged. I don't know what the meaning of the question is. And if that is the objective, I don't see how we can succeed.

Paulson: Is neuroscience the most important discipline for trying to understand consciousness?

Schiff: I think it depends on how one defines neuroscience. Neuroscience is a very broad topic; among the people I know and work closely with are neuroscientists who are also physicists or engineers, or are trained in other fields. So the simple answer is “yes”; if we believe consciousness is a brain process, which we do, understanding consciousness will be centered on neuroscience. But neuroscience *per se* encompasses a large set of activities.

Paulson: Let's come back to this question of subjective experience. Yes, neuroscience can map lots of things that are happening in the brain—for example, *these* parts of the brain have to fire for *this* particular mental state to happen—but is this really getting at what subjective experience is about?

Chalmers: Maybe the question isn't, Is neuroscience the most important thing, but rather is neuroscience *all* you need to explain consciousness? And I think for all the reasons that Danny and others have been saying that it looks like neuroscience alone isn't going to tell us why there is subjective/conscious activity, because there's a potentially unbridgeable gap.

So my view has always been you've got to gather the data from neuroscience, which will be a huge part of the story, but also gather what we might think of as subjective data about consciousness, measured either from the first person point of view or from the sympathetic third person point of view—as when we talk to people and ask them what they're conscious of—and build a multilevel picture that takes seriously the neuroscience but also takes seriously the deliverances of subjective experience.

Kahneman: It's not only neuroscience; real experimental psychology has a lot to say about this. In fact some of the more interesting data are coming from experimental psychologists because they are focusing on the issue—accepting the sort of naive and obvious definition of what consciousness is—that there is an enormous amount of mental activity going on *outside* of consciousness.

There are discoveries being made by experimental psychologists that raise the question, for some psychologists, of what consciousness is for because they don't find anything that cannot be done without it. And I think that this question—what is consciousness for?—is actually being taken seriously.

So, we have the feeling that consciousness is very important for deeper mental activity, for more orderly mental activity, for rule following, but there does seem to be an awful lot of extremely sophisticated stuff that can be produced without it.

Schiff: And that's what makes consciousness-specific measurements very difficult.

Kahneman: That's right.

Paulson: I want to come back to this question about subjective experience. Yes, we can ask people what they're thinking, what they're feeling; we can hook them up to an fMRI and ask them some of these questions and monitor what's going on in the brain. Is that relevant to understanding the essence of what they're feeling? Can science really speak to this issue?

Chalmers: Well, let's distinguish between *gathering* data about what someone is conscious of and *explaining* the data. I can find out what you're conscious of by asking you. This does raise philosophical questions, for example, how can I be sure that you're conscious? Maybe you're a zombie, and so on. But it seems reasonable, under natural assumptions at least, to take what you're saying as a guide to your consciousness; and thereby I can find out about other people's consciousness. But it's another thing to explain this.

There has been a big “neuroscience of consciousness” developing, especially over the last 20 years. And while this area has made significant advances it's still a science of correlation; people draw a diagram of the visual system and *these* bits seem to connect more closely to *these* kinds of conscious states, and so on. But it's still a science of correlation. What we're lacking is explanation: Why is it that all these processes in the brain are in play?

Kahneman: And we have no idea what it would look like if we met it . . .

Paulson: Can we ever get that? Can we ever get an explanation for why these subjective experiences come up?

Schiff: I would guess the answer is yes. But this doesn't get us any closer to it.

Kahneman: I want to raise a difficulty. One thing that troubles me is something that is going to happen from robotics. Some day we're going to have robots with facial expressions that seem to express emotions. And I believe that we will respond emotionally to robots that talk and whose voices indicate emotion; they're going to make sense to us. These robots will look conscious to us. I have no doubt that this is something that is going to happen *before* we understand consciousness; we will have robots that will appear conscious to us.

Paulson: Will appear conscious. Does that mean they are conscious?

Kahneman: I don't know if they are or not; how would we know? It's a judgment that we make about another person. I know my own subjectivity; I believe you're conscious. But my belief about your consciousness, I think, could be simulated by my belief in a robot's consciousness. And where that goes, I have no idea.

Schiff: So let me ask you [*Kahneman*], and maybe David, a question. Do you think that a robot could use natural language—when are we going to have a robot that can adequately use a natural language? I think that's the harder problem.

Chalmers: Yes, that's a hard problem. But going back to the earlier issue, I'll be convinced that the robot is conscious when a robot says to me, "Boy, I know deep down that I'm a set of silicon circuits, but I just can't explain this experience that I'm having of subjectivity." [*Audience and panelist laughter*]

Kahneman: I don't think you can define consciousness by being a philosopher . . . [*Audience and panelist laughter*]

Chalmers: I didn't say this was a necessary condition, just a sufficient condition . . .

Paulson: Let's pursue this question of computer consciousness because, certainly, a lot of people speculate on it. Does a computer have to—does the makeup of a computer have to—mimic the human brain in some way to be conscious, or can its makeup be entirely different from a human brain? What does it take for a computer to start to develop what would seem to be consciousness?

Santos: We have no idea. Computers often trick us into thinking they are conscious . . . Take the Siri function on the iPhone; sometimes it can make you think it's conscious by providing information that seems to require consciousness, "Ooh, you knew there was a Rite Aid there?" But it's a wholly different question to ask whether or not the actual gears of the computer are producing something like a subjective state . . .

Kahneman: . . . and I'm really not sure that we can tell; the third-person requirement for interpretation seems to be overwhelming.

Chalmers: [*to Kahneman*] How do you feel about other people?

Kahneman: Oh, I'm quite sure that everybody in this room is conscious.

Chalmers: Why so sure?

Kahneman: If there really were convincing robots sitting in this audience, I would not be able to say, "Oh, *that* one is not conscious." The evidence that I have of the consciousness of other people can be produced artificially.

And if we ask, “Can robots speak a natural language?” . . . —I had a teacher in philosophy many years ago, Professor Bar Hillel; and he was asked sometime in the 1950s about when computers would understand language and he said, “Oh, never, and by ‘never,’ I mean the next 15 years.” So, we really don’t know what “never” means in this context.

Paulson: [*To Schiff*] Did you want to follow up?

Schiff: Well, I think that part of the problem is our best evidence for consciousness in anyone else is our own subjective experience. And in essence, this brings up the issue of how little we get out of external observation to help us understand our intuitions about consciousness and what it could be.

In my very bottom-up approach to this, the question that I and my colleagues are dealing with boils down to: Could this system (person) recover consciousness? Such a question is immediately mechanistic and within the context of the human brain; it is very challenging and requires consideration of additional questions such as, what brain state can produce consciousness? I think we’re getting much better, even in causal efforts, to get to the brain state question. But the issue of subjectivity and how that happens mechanistically, I think, is opaque, as you say [*to Chalmers*].

Paulson: Laurie, I want to pursue this in terms of animal consciousness. You study cognition in primates, in particular. What kind of comparisons can we make between human consciousness and nonhuman primate consciousness?

Santos: This is a difficult issue, to ask about animal subjective experience. For many animals, we can’t be with them without thinking they have subjective experience. [*To the audience*] To those of you who have a pet dog, I bet you implicitly assume that the dog has a deep and rich subjective experience, even though you just don’t know. Even if we ignore the subjective experience question and ask, What are animals thinking?—which I believe we have better measures to gain some traction on—it is still puzzling.

I think the more we get to know about animals and the more fascinating things they do—they’re not using natural language, but they’re making incredibly complicated decisions, incredibly complicated evaluations—they’ll be seen as having preferences and other behavioral signatures that we associate with subjective experience. But again, even though we often think animals have lots of subjective experiences—we don’t really know if animals do in fact have them.

Kahneman: My sense is slightly different. I think that emotion is very important in our attribution of consciousness to other people and to animals. Computers can compute very complicated things but this ability is not consciousness.

Instead, it is the emotional connection to someone (or animal) that gives us the intuition that it is conscious. This is an interesting psychological question: What makes us *feel* that something is conscious besides ourselves? The answer will include psychology. And whether that psychology can sustain a science, I’m very skeptical. The example of the computer or robot fooling me by looking emotional or not fooling me by having true emotion just isn’t compelling to me. That’s why I never got quite caught up in this issue.

Chalmers: One of the classic philosophical problems, the problem of other minds—How do you know that anybody has a mind? How do you know who or what has a mind?—is cropping up practically in some ways within the science of consciousness. How do we know that animals are conscious? How do we know that computers are or are not conscious? And in Niko’s work, how do we know that people coming out of coma and some vegetative states are conscious?

What we find is that people are very imaginative and creative, and there are techniques that are being developed that, while they don’t solve the philosophical problem, are criteria for consciousness that seem to fit with our normal practices of ascribing consciousness to people in everyday life and elsewhere. There is beginning to be a field of what we might call “the psychology of other minds,” which is what

Danny is alluding to, in which the aim is to determine the criteria that ordinary people use for ascribing consciousness. It turns out the criteria for consciousness seem to include things like pain and emotion, and so on. Simply being a thinking thing, without emotion, does not correlate as well with consciousness.

Schiff: The question of whether somebody is conscious can come up in ways that are just unimaginable until one is actually faced with certain patients at the bedside. This comes up in almost all of our work. And there are some cases that I still go home every night and I think about and worry a lot about because we know that some people can be locked in. Anybody who's seen the movie *The Diving Bell and the Butterfly* has encountered an example of this; somebody's fully conscious yet they've lost their motor function; they are "locked in."

Paulson: So from an outside perspective such people seem totally without consciousness.

Schiff: No, not exactly. If you're a good examiner, you can figure out that they're conscious right away; it's no problem.

Chalmers: How do you do that; how do you know they're conscious?

Schiff: Because you have a reliable communication channel.

It's probably worth describing operationally what happens when looking at patients in coma versus a vegetative state. And at least by definition, when a neurologist looks at a patient and identifies him/her as in either coma or vegetative it means the same thing from a behavioral point of view: there's no evidence of responding to the world, no evidence of taking in sensory information. In contrast, the difference between coma and vegetative state is a technical one and has to do with function within the arousal systems in the brain stem returning and producing a changing pattern of eyes open/eyes closed periods during the day. This is not related to sleep and wake; and it's not associated with the kind of electrical activity you see in sleep. It's a change associated with the typical recovery pattern after coma (with the caveat for the neurologist and the audience that there are occasionally eyes-open comas associated with a particular kind of injury).

But at the very border of vegetative and conscious states is the next level of recovery, which is now being called *minimally conscious state*, where we begin to see unambiguous signs of response to the environment. In this sort of gray zone between these two conditions (conscious and vegetative) one finds simple responses such as tracking of a visual image or eyes looking toward a sound. Although these responses don't seem very different than opening and closing eyes, it turns out that they're increasingly recognized as signs linked to huge differences in the prognosis of patients.

Recognizing that such small changes have important prognostic consequences for patients is not being dealt with very well in the medical community for a number of reasons, including finding room for patients in minimally conscious states who may go on in that state for weeks and could recover, even to the point of ambulating and walking around in a year, but yet are not getting adequate therapy. So, this is a major issue.

And as patients recover more function, they might begin to respond to commands. This level of consciousness is obvious. Operationally, once we can communicate with somebody with a reliable "yes" or a "no," then we can assume consciousness.

Chalmers: But there are locked-in patients who, as in *The Diving Bell and the Butterfly*, only have control over their eye blinking.

Schiff: Right, eye blinking or eye or head movement.

Chalmers: The question that obsesses me about these patients is, how do we know there aren't locked-in patients who *don't* have control of their eye blinking?

Schiff: Right. So you anticipate why I was setting all this up for the audience. That's exactly right. So, two examples. One is a published case of a patient who was locked-in classically with a brain stem injury and had an unusual extension of the injury into the auditory system. The auditory system is usually very well preserved because it goes to both sides of the brain. But this person ended up with a central auditory agnosia, which means the patient could hear but couldn't really put together complex sounds; the condition was not recognized for many months because the patient relied to an extent on lip reading. However, during tests of the patient's cognitive level, the lip reading would fail and it seemed like the patient inexplicably fell off the curve and was, consequently, judged to be minimally conscious or cognitively impaired. It wasn't until a colleague of mine finally decided to draw and write questions for the patient that it became clear that he was fully conscious and just needed to have the visual representation of the words. So in this case, the patient was locked-in and conscious.

The second kind of case is far more troubling. These are patients who at times seem just like a locked-in patient, they look down for "yes," look to the side for "no" but only have accurate communication for, say, 2 hours a day; they attempt to work with a brain-computer interface but aren't successful. These patients seem locked-in but they have ambiguous signs of consciousness. The challenge is that such patients bring a sense of urgency to treat in the hope that they will recover, yet it's not clear if this is possible, since one cannot decide between two possibilities: Is it that they're not able to control their motor function, or is their conscious state having a problem? How would we be able to make meaningful measurements to distinguish between these?

Paulson: This is so fascinating. . .

I want to pick up the thread that we were talking about earlier about animal consciousness because there are many fascinating questions here. The question that so many of us ask is, what is it that makes us human? Which leads to additional questions: What is it about human consciousness that sets us apart from the rest of the animal world? Is it something fundamentally different about us or is it just a matter of degree? Are chimpanzees basically like us but at a lesser level?

Laurie, what's your sense?

Santos: As I mentioned earlier, in terms of subjective experience, we don't have ways to measure what chimpanzees experience. In terms of cognition—in terms of how animals think—we're starting to get some important hints about what makes humans different.

I bet if I did a show of hands, most people would think that language makes us unique. And while I think language helps, because it's hard to express oneself without language, my view is that language is actually a red herring. I'm kind of in the minority of animal researchers who think this. Consider, for example, if pigeons had language; I believe they wouldn't really have anything interesting to say [*audience laughter*].

And so lately people have started to think that, actually, one of the things that makes us human is the kind of thing that we're engaged in here right now, which is not just that we're communicating with natural language but also that we're motivated to share what we're thinking. That's the power that language provides: it's built on a motivation to share. And we're getting new hints that it seems like other animals, particularly even our closest relatives, chimpanzees and bonobos, don't seem to have the same kind of motivation to share what they know about the world. This predicts some big differences and it means that other animals are not going to communicate in the same way as us; they're not going to innovate and share what they know.

Paulson: Let me mention something that Jane Goodall once told me during an interview. She's fascinated by the whole idea of thinking without words, thinking without language. And she said that if she could spend just a few minutes inside the mind of a chimpanzee, she would learn more about what being a chimpanzee is about than all the decades of research that she's done. Does that resonate with you?

Santos: Oh yes. If I could take some sort of, I don't know, not time machine but a sort of mind machine and dive in to an animal's mind it would be fantastically interesting. But I wouldn't pick a chimpanzee;

I'd pick something much further away on the evolutionary tree, an ant, for example. People always jump to insects when they're trying to pick something that isn't conscious. But the fact is, sometimes insects are doing some amazingly complicated things that mimic what humans do in rich ways.

My favorite example is from E.O. Wilson, who talks about an ant's reaction to death. In an ant colony, one of the things ants have to deal with is debris and death; it turns out that if you put a dead ant inside an ant colony, the ants will take it out, identify it, take it to an area of refuse and lay it in what folks call an "ant cemetery." This behavior led some animal cognition folks to say, "Wow, maybe ants have a concept of death; what are they using to decide that another ant isn't conscious?" It turns out, however, that this ant problem is actually pretty straightforward; the behavior of ants—"taking care of their dead"—is due to a small chemical called oleic acid that ants emit from their exoskeletons when they die. In fact, if oleic acid is placed on living ants that are running around, other ants in the colony will grab them and throw them out of the colony. [*Audience laughter*]

Paulson: Let me follow up on this. I'm going to ask you all to speculate on something. How far down the animal chain do you think consciousness goes?

Santos: Well, first of all, I would say "across" the animal tree rather than down the animal chain. I was a college student with the late Don Griffin, who was a scholar of animal consciousness, and he made the claim that since we don't know of anything different about the organic matter that makes up nonhuman brains—nothing that makes it different than human brains—the most parsimonious answer is consciousness is likely to be seen in all other organic creatures that we call life. So that's one view.

Kahneman: In the absence of a criterion how do we know? The ant story you tell is very interesting because it ties up with what we were talking earlier about: emotional responses. Here's something (ant behavior) that looks like an emotional response and our intuition is that we can empathize with ants; but then it turns out the behavior is due to a reaction to a chemical, and then we want to say, "Forget it, there's no evidence of consciousness."

Well, this argument is flawed every possible way. On the one hand, we do feel that understanding the chemical cause ruins our intuition. But on the other hand, maybe what we're responding to when we attribute consciousness to other humans is the equivalent of a chemical for ants; for us it's just an emotional expression that we can empathize with.

Chalmers: It's very hard to find the place where consciousness "gives out" as you move across the natural order. If one looks for a capacity without which there wouldn't be consciousness, for example, language or reasoning or certain kinds of emotion . . .

Paulson: . . . do you need a brain to have consciousness?

Chalmers: Well, this is one of the questions, Is consciousness restricted to animals with brains? We've got consensus on the biological on animals. But in fact you [*to Santos*] mentioned Don Griffin. His brother, David Griffin, is a philosopher who has come out in favor of panpsychism—the view that everything has a mind; that there's some element of consciousness at the very bottom level of the natural order.

Paulson: Dave, haven't you come out with something similar? We have talked about this. You have suggested that consciousness cannot be reduced to physics; that in fact consciousness may be a property of nature in itself.

Chalmers: We have things in science, like space and time and mass, that cannot be fully explained in terms of things simpler than themselves. This is not unscientific, it's just something we're used to in physics. Outside of physics, we're used to explaining things in terms of other, simpler things. But this is a strategy that doesn't seem to work well for consciousness.

So I have been led to speculate that we should take seriously the idea that consciousness is fundamental. Once you do that, it's also natural to speculate—and it is just speculation—that consciousness may be present at a very fundamental level of the physical natural order and that, for example, David Griffin may be correct, that consciousness is to be found—some element of consciousness is to be found—in fundamental particles.

Paulson: What *is* that element of consciousness, then? Is it made of something?

Chalmers: Some people prefer to say “protoconsciousness.” And the panpsychists tend to say some “precursor to consciousness” . . .

But in fact we don't understand this. We don't understand the nature of matter, and we don't understand the nature of consciousness. When it comes to consciousness, we're in the dark. And yet philosophers speculate for a living; we try to describe what kind of picture of the natural order will make sense of consciousness [*audience laughter*].

Paulson: Okay, Danny, you're shaking your head.

Kahneman: I don't want to spoil the fun . . . [*audience laughter*]. But earlier I think we were agreeing that what we have are intuitions about consciousness. And when we talk about consciousness as a noun rather than about the intuitions about consciousness, there's virtually no limit to what we could argue, because we actually don't know. All we have are intuitions about consciousness. So, it is the legitimacy of that question—What/where is consciousness?—that I would like to question. If we don't know what it is, if all we have are our intuitions, then ultimately all we can do is study the psychology of these intuitions. And that is a very different thing than studying the ontology of consciousness itself.

Chalmers: So we don't know where consciousness is; we certainly don't know that consciousness is present in nonbiological systems. At the same time we don't know that it's not present. We simply don't have data about that.

Kahneman: But I mean if we will never know—that is, all we have are our intuitions about the presence of consciousness—then I genuinely do not understand the question. I don't understand what we don't know.

Chalmers: I think what we need to do is to build a theory of consciousness that explains the data that we know about, which is human data . . .

Kahneman: . . . that explains our intuitions? . . .

Chalmers: . . . that explains our first-person intuitions about consciousness.

Paulson: But doesn't this come back to the question of, Can science explain consciousness? The idea that's on the table here, it seems to me, is that maybe there are some dimensions of consciousness that are beyond the explanatory power of science. Is this a legitimate hypothesis to entertain?

Schiff: Well, it's not a very interesting one . . . [*audience laughter*]. Because it doesn't do any work for us. I would say that we should just bracket that possibility and work as if we could discover enough information about how certain things work mechanistically to gain an intuition that's precise enough about how things we consider to be conscious happen in the human brain. The question of where does consciousness start on a phylogenetic spectrum is a hard one; without a mechanism attached to what we already are trying to solve about whether we're conscious or not, it's not a meaningful question to me. I don't think a jellyfish is conscious, at least in a way that makes any contact with me or does any work for me to help me understand

the problem of assigning consciousness as a mechanistic possibility for a physical system, which is a brain, usually, when I'm thinking about the physical system.

Kahneman: But if it is true that we could be fooled into thinking that a robot is conscious—a fully expressive, emotional robot—this really changes the picture entirely. In this case all that is left, I think, is the psychology, because asking the ontological question, Is the robot *really* conscious?, I'm really not sure that there is anymore that we can do besides to say that, "Yes, we think it is; we feel it is."

Santos: I think that studying and figuring out our intuitions about what is conscious is an important question for two reasons. One is that our intuitions are often wrong. Almost everything we know about our intuitions suggests they're fantastically wrong. However, they govern a lot of our behavior and a lot of our judgments about things. Take an example in politics; we currently have questions about whether corporations should have rights; our intuitions about whether corporations have subjective experiences probably tell us something about what we should be doing to them. And another example. For some of us here tonight there's a question to confront when we go to the reception: Should we eat things that have meat in them from certain animals? Our intuitions about whether we should are likely governed by our intuitions about whether or not those animals have conscious states.

So, I think understanding what our intuitions are telling us is going to be really important and meaningful, even though we may find out those intuitions are in fact wrong.

Paulson: Dave, I want to come back to you. Do you agree with what the rest of the panel, or some of the panelists, is saying, that basically some of these questions are not relevant because science has no handle on some of the larger philosophical questions about consciousness?

Chalmers: I'm not saying that Niko should give up his day job . . . But I do think it is a meaningful question whether jellyfish are conscious, whether or not the discussion goes further. And there is a fact of the matter about it, which maybe we're not in a position now to understand or describe—maybe not ever. However, there may be ways eventually to get at this through scientific methods, if indirectly.

Here's what I think we have to do. We have to start with the cases of consciousness we know about, the cases where we have data—roughly, the human case. Build a theory, an explanatory theory, that connects consciousness to, for example, brain processes. I think of this as trying to make an abstraction of the fundamental principles that connect brain process to consciousness. It could turn out that the most successful theory that explains the data we have says that consciousness is generated by certain kinds of complex processes, certain kinds of reasoning, or certain kinds of complex recurrent structures, and so on. We would then be in a position to extend the theory to other cases. Where that structure/complex process is not present we should not expect consciousness to be present. It could turn out, on the other hand, that the theory of consciousness that best explains it ties consciousness to some other basic properties, for example, to information and information processing in the brain. Similarly, we could extend this theory to other cases; and it may be speculative because we can't measure consciousness directly in those systems.

Paulson: This might help to explain computer consciousness then. If consciousness ultimately is about information, a computer might have an integrated information system as the human brain does.

Kahneman: Well, that could be; but this would lead, I think, to a conflict in intuition.

[*To Schiff*] It could be that your research could isolate an area in the brain that is associated with consciousness. I'm speculating. And then if we found animals that don't have that area of the brain, we would have some reason to say they don't have consciousness. But that would really conflict with the intuitions associated with a robot that certainly doesn't have those areas of the brain but can generate in us the intuition that it has consciousness. So, I don't really see a way out of this.

Schiff: The problem with the robot as an analogy, I think, is that it typically isn't the case that most intuitions about consciousness come from observing other people; importantly, they also come from introspective aspect of examining natural language and having—or at least taking the attitude that we have—subjective experience, and then acting accordingly and sharing it, as you would say.

Kahneman: But you could surely program a robot to do all of the above . . .

Schiff: You could; but if you knew it was a robot, then that might be the reason why you wouldn't attribute consciousness to it; and that wouldn't be a bad reason. Unless we have a mechanistic understanding of consciousness and then instantiate that mechanism in the robot, we wouldn't likely attribute consciousness.

Chalmers: So, Danny, how about we take your neurons and replace them one at a time by silicon chips that are functionally isomorphic to the original neurons so that half your brain is silicon, and we ask you if you are conscious . . .

Kahneman: There has been a lot of speculation about whether a conscious brain has to be made of “meat” and I don't see any reason why it would have to be. So, if you found silicone substitutes and the functioning remains the same and the emotional expressions remain the same, I see no reason not to attribute consciousness.

Schiff: [*To Kahneman*] That transfers the problem the way David is talking about because, as you say, it could be information or it could be some other aspect or property of complex matter—the brain is just one example of having that particular property—that's essential. And once we understand that aspect or property, certain things will become transparent about how consciousness happens.

Chalmers: I do think it's more likely to be information than biology . . .

Schiff: Right, like plasma physics; or something about condensed matter . . .

Kahneman: If it is information, then that doesn't fit with our intuitions, which are driven primarily by emotions; our attributions of consciousness are driven by emotions. When we think about consciousness and think about information processing, there is really a deep disconnect between those two.

Paulson: Let me ask you all about the research project going on now, what some people have called the *connectome*, the extremely ambitious project to map the neural circuitry of the brain. It's almost beyond comprehension to think about it: there are nearly 100 billion neurons in the brain, and trillions of synapses. Let's just say theoretically somewhere down the road, this can be mapped. How close are we to understanding consciousness then?

Schiff: I know a lot of the people on that project and so I feel like if I answer, I'm going to be sharing one of their ideas, so maybe I'll attribute it. One of my colleagues, Partha Mitra, who works out at Cold Spring Harbor, is trying to develop a stitched digital atlas of the detailed connections of a mouse connectome, which is one of the projects in this space. Mitra points out that for many cells their connections are so manifold across the brain that an understanding of a full map is going to require different theoretical models and laws about how systems that have such an architecture might even possibly work.

So, I don't know that we're closer to understanding consciousness, but we're going in the right direction.

Paulson: Dave, what do you think?

Chalmers: An absolutely fascinating project. I actually had lunch today with someone who's working on this project, the brain activity map. They're trying to get funding to produce a map of a whole brain beginning with something small, say a fly, and then in 50 years a human brain if the right kind of imaging techniques can be developed.

I think that once we've got that kind of tool—a human brain activity map—neuroscience is obviously going to be revolutionized, because in neuroscience, we're at the mercy of our tools and the tools today are very limited. Yes, new imaging techniques come along, like fMRI, and suddenly experiments are transformed; but it's arbitrary what the new techniques give us access to. In contrast, having access to every neural firing in the brain and every connection is going to suddenly put us in a position where the mechanisms start to become transparent to us. But these will still only be the mechanisms. I think what we end up having is a situation where we get an extraordinarily sophisticated science of the correlations—presumably, we'll be able to manipulate the brain and simulate certain things; one will even be able to do it in principle to one's self. Actually, there was a psychologist, Paul Meehl, who wrote an article back in the 1950s called “The Complete Autocerebroscopist.” It was exactly this scenario in a thought experiment: you have a picture of your brain; you are the experimenter; you are in this scatter of the complete data about your brain in front of you; you can experiment on your own brain and see how your experience changes.

So, in principle, with a brain connectivity map we would have a vast trove of objective data about the brain, and with our own introspection we have a vast trove of subjective data about consciousness. One might think it would then be possible to abstract out the relevant kinds of principles that connect the objective to the subjective. I don't think this would mean that we would bridge the mind/brain gap, but we would have boiled down this conundrum to the simplest possible principles.

Kahneman: Would that not be correlational?

Chalmers: I think the best we can get is correlational, but we can get better and more systematic correlations . . .

Santos: . . . we can get lots of correlations.

Chalmers: Well, what we have in physics ultimately are some fundamental principles that are in a certain sense correlational, like the law of gravity; but in the end we get a simple principle that generates the data. I think they say in physics that the ultimate goal is a set of laws so simple you can write them on the front of a t-shirt. If we got to the point where we had laws of consciousness connecting physical processing to consciousness so simple one could write them on the front of a t-shirt, maybe that wouldn't remove the mind/brain gap (because we'd still have the abstraction), but I think we would call that a pretty powerful theory.

Paulson: Does one need consciousness to have a sense of self?

Santos: Well, since we can't fully measure consciousness, it's tricky to ask . . .

Kahneman: . . . and sense of self as well . . .

Santos: . . . and sense of self is also very difficult to measure. There's a long history of work in the field of animal psychology that attempts to measure a sense of self, and a lot of that history, I think, points to some creativity on the part of researchers but, ultimately, not great methods despite the fact that a lot of creativity was employed.

Paulson: The working assumption in this area is that if the animal can recognize itself in a mirror, then it has a sense of self.

Santos: Yes. And in fact even today we find researchers testing new animals with new mirrors and new kinds of marks on their forehead trying to see if they recognize themselves. And the pattern in the experimental data seems to be that the standard big creatures who seem smart like us seem to do it pretty well, so chimpanzees and elephants do well; but other critters who are smaller and don't look a lot like humans don't seem to do it.

[*To the audience*] Consider your sense of self and your set of preferences and your future goals. If I tried to reduce the presence of that subjective experience and information to the simple act of you noticing a splotch or marking on your face when you walked by the bathroom mirror, you might feel that measure was kind of missing something. And I think that's the general sense in animal cognition right now, that mirror self recognition can tell us something about which creatures can recognize themselves in a mirror, but it's not capturing a meaningful sense of self.

This is the big puzzle and challenge for people who work on animal cognition; we have good ideas of the kinds of things we humans do, but capturing and measuring those same abilities in an animal that lacks a natural language is really tricky.

Paulson: I think it's time to go to the audience here.

Audience member 1: I would very much appreciate each of the panelists giving a definition of consciousness.

Schiff: Good question . . .

Chalmers: I think it's very hard to define *consciousness* in terms of anything more basic than consciousness, just as it's very hard to define *time* and *space* in terms of anything more basic than time and space.

But there are things we can which at least I think are helpful. There's a phrase due to Thomas Nagel, who was mentioned earlier, who wrote the article "What Is It Like to Be a Bat?" You might say that a system is conscious when there is something it's like to be that system—so it's something it's like to be me; it's something it's like to be you. But importantly, assuming you're not a panpsychist, you would say there's nothing it's like to be that [*points to a cup on the table*] cup. So, likewise, a mental state like seeing will be conscious if there's something it's like to be in that state; for example, there's something it's like for me to see you right now, but there's nothing it's like for me to do some computation in my cerebellum.

So, I don't know if that's a definition exactly, but it's at least a way of getting a grip on what I'm talking about when I talk about consciousness.

Paulson: Niko?

Schiff: Credit to William James, as I think most neurologists use some variation on James' definition of consciousness, which is that consciousness is awareness of the self and/or the environment . . .

Chalmers: . . . you just used the word *awareness* . . .

Schiff: [*smiling*] . . . right, which is tautological, as David, the philosophy professor, quite rightly points out; although I think James got away with it!

Santos: If the philosopher can't give a definition of consciousness that will satisfy you, I'm surely not going to [*audience laughter*].

Kahneman: Well, I'm going to question whether there is a satisfactory definition, because I think all we know is what intuitions we have about what is consciousness. And if that's all we know, then defining consciousness as if it exists independently of our intuitions is an exercise that I have no idea how to conduct. I know that very intelligent people deal with this and spend their lives doing it, but it is an exercise, frankly, that I have never understood; that's the frustration I've been expressing.

Chalmers: Actually my favorite definition is: *consciousness*—that annoying time between naps [*audience and participant laughter*].

Paulson: [*To audience member 1*] Excellent question.

Kahneman: We can go with that . . .

Paulson: Next question.

Audience member 2: I'm interested in the difference between *kinds* and *degrees* of consciousness. With the vagaries of evolution and the variations it throws up, why should we think of capabilities as being the only things that vary in consciousness? Why couldn't it be like the examples given, namely, gravity and electricity. These are very different fundamentally, yet they're both forces. Similarly, why couldn't there be different kinds of consciousness in terms other than in capabilities?

Paulson: A response?

Santos: I think once we get a good metric for a subjective experience, once we know how to measure it in us and other beings, I can't help but imagine that we're going to find differences in subjective experience across species. We don't yet have a great way to measure consciousness in me and we definitely don't have a way to measure consciousness in a jellyfish. And yet while there's something that it feels like for me to feel embarrassed—there's that state that I experience—I would bet there's no similar thing that it's like for a jellyfish to feel embarrassed. My guess is we're going to find gradations of consciousness once we actually have a good way to measure it. But the problem, again, is that we don't have a great way to measure it now.

Paulson: Let's take another question.

Audience member 3: I was wondering about using a scientific method that's based on objectivity to study subjectivity, and whether in fact part of discovering and understanding subjectivity in consciousness requires a new scientific method.

Paulson: Good question. Do we need a new science here?

Chalmers: I think there can be objective facts about subjectivity. I think it's an objective fact that I'm conscious; it's probably an objective fact you're conscious. We're faced with this problem in science all the time. Can I know for sure that the external world exists? We have to make certain assumptions to get that off the ground, for example, perception is one guide to reality, but not an exhaustive one.

Paulson: But do we need a new scientific frame? Do we need a new scientific paradigm here to crack the consciousness problem?

Chalmers: I think we have to do at least this much. We have to take the data, subjective first-person data, *seriously*, because those are the data we have. Why do we believe we're conscious? Because we experience it.

There have been moments in the history of science when these issues have been taboo and off-limits, when everything had to be studied and measured from a third-person point of view.

I think consciousness is something we know about from the first-person point of view, and so the science of consciousness has to admit observations from the first-person point of view as data, because I think such observations are data.

Kahneman: I'm not sure anybody is debating this [*Santos and Schiff nodding in agreement*]. First-person accounts are accepted as data; there's no problem. The question is what do those data tell us? And can they lead us to a solution? And that's quite unclear.

Paulson: Let's go to the gentleman right over there.

Audience member 4: My question would be, let's say we can define and all agree on what consciousness is, and let's say we have all the techniques and all the science to identify what is a conscious state, what is an unconscious state, whether a robot is conscious or a jellyfish is conscious, etc., what do we do with that information? What is the purpose, what is the application? Maybe thinking that way could help us back-engineer how to look at consciousness and what it is?

Schiff: I think there's an ethical obligation to understand this, and there's a humanitarian issue behind a lot of what neurologists want to know about measuring consciousness. For example, if somebody is conscious and yet can't express him/herself or can't communicate with family members, and we can figure out by using a science of consciousness a way to enhance their capacity to do these things or to give back a level of activity by doing something intelligent, that's an inherent good, at least from my point of view. That's one of the motivations.

Paulson: Isn't there a profound ethical question, which you are alluding to, here? Take the example of someone in a horrible car crash; the doctor, the people on the scene say, "Sorry, this person is brain dead now." Do we pull the plug?

Schiff: I can make that problem harder and harder and harder. In medical settings it keeps getting harder. At times we are stuck with ignorance and situations in which, without a good model, mistakes are made. In many cases, all one can do is give the best information available to families and let them make decisions within the range of uncertainty. This is becoming more and more uncertain as the science is evolving, not less.

Paulson: So isn't the implication then that we should be much more reticent about terminating life?

Schiff: But the flipside of this is that you also don't want to commit people to indefinite vigils for patients who are not going to recover. At best reasonable assessments and average answers are given because there are no definite models. For the reasons that we've heard today, most of what is done in consciousness science applied to medicine is statistical. Somebody's in a vegetative state; they had a type of injury; there aren't a lot of other measurements that can be done if they're in a zone where it's not certain. Yes, sometimes things are simple. Sometimes the brain is dead; brain death is death—it's a diagnosis that doesn't have a prognosis. But if the person is in minimally conscious state a month after a very severe injury, things become harder. One might be able to statistically estimate a level of function to be at X, Y, or Z, but then it's an issue of value, an issue of what contact person is going to have with their family. Often, the definite answer to such needs is completely unknown. There's a range of what is acceptable and there is a range of what individuals will see as meaningful human contact with somebody they love and have known for their whole life.

This happens in Alzheimer's disease all the time, right? This is something everybody's familiar with: an elderly person who's slowly slipping down out of contact. If we had some way to bring them back so they could talk to their grandchildren for a year, this would be a very important thing. And while it may not save the person's life and might not mean that they wouldn't slip out of contact at some point in the future, it would be important.

This is the kind of thing that a science of consciousness will make more law-like and predictable, and allow us to talk more intelligently about such questions as, What can and can't happen for my loved one?

What we should do; should we do this; is this something you want to do? Such conversations are never going to be easy, but they could be better informed.

Paulson: Next question.

Audience member 5: Thanks for an awesome panel. It's curious to me that nobody in the panel is representing a thoroughgoing reductionist account of consciousness of the type that, say, Daniel Dennett argues for in his book *Consciousness Explained Away*.

There is a fundamental problem that nobody's really addressed, which is that we're using the brain to study the brain, so there's a kind of tautology here that we're locked into. And I think Holding, the biologist, once summarized it very well; I can only paraphrase it but basically he refused to believe that his mind was simply a biochemical process because otherwise he would be sawing off the branch on which he sat and would have no reason to trust his science at all.

So, if consciousness is purely a biochemical process, which nobody on this panel has actually represented, and there's not something emergent or transcendent in the difference between brain and mind, why should we even believe scientists when they speculate on these issues?

Paulson: Who wants to take that? [*Panelist laughter*] The philosopher? David?

Chalmers: Well, I think it was directed at a reductionist, which I am not.

But I think the brain can study the brain. Microscopes can study microscopes and so on. There's not a rule in principle against systems directing scientific inquiry at themselves.

The question, I suppose, really is, how do you even know the brain is there? How do you know you're not in the Matrix? Descartes said it's all produced in our minds by a Cartesian evil demon who wants to fool us into thinking that we're scientists doing experiments on people's brains, when in fact it's all a big fantasy.

I think in science you basically just have to assume that the deliverance is a perception or a guide to reality, because without doing that you can't start doing science. And I think it's the same with the science of consciousness.

Kahneman: There are paradoxes where the best response is simply to walk away . . . [*Audience laughter*] You're not going to solve them.

Paulson: When I interviewed Dan Dennett and asked him, "So how big a mystery is consciousness?" He said, "It's not a mystery; there are just lots of puzzles."

Next question.

Audience member 6: Thus far most of the discussion has been about people who are scientists; but there is a whole bunch of other people out there, not the least of which are great writers. On the panel, who would you consider, as a writer, to have made a contribution to crystallize for you what consciousness is? There are some who come to my mind—James Joyce, Oliver Sacks, and maybe Julian James and the breakdown of the bicameral mind. I mean, there are a great many people who are really concerned with this.

Who do you like as a writer who may have contributed to your discipline?

Schiff: That's a great question.

Santos: [*To Paulson*] You mentioned one of the ones that I would put out there, Jane Goodall. I think good writers that tell us something about consciousness are just really good at doing the game of "what it would be like to be a ___." And in my world, knowing what it's like to be a chimpanzee helps you hang out with a lot of chimpanzees.

I think all the time Goodall spent trying to get inside their mind is very helpful. If you read some of her pieces aimed at a popular audience to try to describe what it's like to be a chimp, for example, the chimp she named David Graybeard or one of her other favorites, you get a fantastic look at that.

To me it's about writers being good at getting inside the heads of other individuals, which is not necessarily a scientific understanding of consciousness or figuring out the problems that we've been up here talking about. It's having good intuitions about being inside somebody's head and then describing this to the readers.

Chalmers: I'm going to go for Proust because I'm really interested in phenomenological methods: representing. One of the big challenges for science is produce tools to accurately describe and represent states of consciousness, which we're not terribly good at. We need a better phenomenological method. And with Proust we see a master phenomenologist characterizing his states of consciousness in vast, gory detail.

I think this is one of the things we can look to writers to do. The question is, can one take this and make science of it?

Paulson: Other literary influences?

Kahneman: All those I would mention have already been mentioned. Oliver Sacks would be one that I'm most current with. Again, the issue arises that these writers speak to our intuitions; they do something that feels right. This is how I've thought about many philosophical questions ever since I was a child; the psychology of our philosophical intuitions is what's fascinating to me. And the psychology of why some accounts of consciousness appear appealing and others don't is fascinating to me, even if I really don't think the problem can be solved at all. But the psychology is good. Or can be good.

Schiff: This is a great question. What's came to mind immediately when you asked about a writer who has influenced my work or thoughts on consciousness and self is Helen Keller, who is an incredibly amazing writer. At the end of her life, Keller wrote a book called *Teacher* about Ann Macy Sullivan, in which she talks about her feelings about who she was, or what she was, before Sullivan taught her language. And it's amazing because she doesn't attribute a self-identity to the conscious being that she was; she actually labels it "phantom," and she describes it in detail as an entity that was just being pushed around and reacting, usually violently and with limbic drives. Keller is such an eloquent writer. And while this is not the focus of her book—it's mostly about how Sullivan taught her and what it meant to her—that section, for me, is very interesting. I've been reading it and thinking about measurements, and thinking about this general issue.

Audience member 6 again: Can we ask the narrator?

Paulson: Me? Williams James is one of my heroes, and has been mentioned. I'm just astonished at how contemporary he still is. A book like *Varieties of Religious Experience*—I have a particular interest in questions about science and religion—in which James asks questions that most people don't even talk about still. I find him stunning.

Let's go to the next question.

Audience member 7: One of the things the panelists talked about was how when people are showing emotions they can be considered conscious, or when people are doing complicated calculations they can be considered conscious.

When I walked through work today, I was probably not aware of any of the steps that I was taking; I was doing it unconsciously, but they were very complicated calculations. So, why would seeing emotions or doing a calculation be necessarily related to consciousness?

Santos: When I mentioned doing calculations as part of our intuition that a creature is conscious, I think I meant doing calculations that are like the ones humans do. As somebody who studies animals, I watch people's intuitions about which creatures have subjective experiences "turn on" when they watch animals doing things like we do. Take the ant case again. I think one of the reasons people are intrigued by the ant activity of "reacting to death" and "taking dead individuals out to a cemetery" is because those are the

things we do. So, certain activities/behaviors can trip our intuitions to lead us to think, “Oh, that creature must be like us.” I think Danny’s right, that acting in emotional ways like we do is a special way to trip up these mechanisms.

I’m not sure anyone here has seen the Ikea television commercial about buying new light bulbs in which they used a Pixar-esque computer generated image of an old lamp that was about to get thrown out; the commercial has very emotional music and the old lamp is sort of hunched over and looks sad when a new lamp comes in; and then the old lamp gets tossed out in the rain and it stands in the rain. Watching this video, one can’t help but empathize with the old lamp’s feelings of rejection. But the punch line comes from the Ikea salesperson who says, “It’s just a lamp; it doesn’t have feelings! Buy a new one!”

The point is we can’t help having intuitions when we watch a lamp that’s exhibiting behavior that looks like human behavior; and so my instinct is the creatures that trip our intuitions easiest are those that have subjective experience most like ours.

Kahneman: You see a dog feeling guilty; that’s a powerful intuition. There is a behavior that makes us empathize with guilt, and dogs can produce that behavior . . . , and are often rewarded for it [*audience laughter*].

Paulson: We’re running very short of time. Let’s go to the next question.

Audience member 8: I think no matter what we all believe about consciousness, somehow we all believe that we are conscious, on some level. And we all believe this so much that we all decide to take a bunch of time out of our day to come and listen to this panel. And we obviously evolved that way. So my question is what are the advantages of being conscious? And why are we? What does it give us?

Santos: I think this question gets even more complicated when we start to realize that many of the decisions we make that *feel* the most conscious aren’t necessarily the ones that have the most impact for us. Danny can speak to this more than I can. But I think the puzzle gets bigger because if you think about all our mental processes that aren’t conscious, like not getting hit by car on the way here tonight, from a natural selection perspective those mental processes could be the ones that are doing a lot of the hard work of keeping you alive, getting mates, and so on.

Chalmers: We just don’t have an evolutionary answer to that question right now. I think if we did, the theory of consciousness would be in much better shape. One way to put the basic question of consciousness is, Why didn’t evolution just produce a race of zombies?—physical systems that went around doing things without consciousness at all. But evolution didn’t; consciousness is there.

So, on the one hand you might think there are reasons for consciousness and that in fact has some important function; but on the other hand, every time someone proposes a function it looks like that in principle everything can be done without consciousness.

Paulson: Next question

Audience member 9: Is consciousness the supreme biological function? You seem to even suggest that it is, at some point, akin to life. Or is it just that we’ve all evolved the ability to think about it so we’re sitting around talking about it as though it is.

Chalmers: I work on consciousness, so of course I’m going to say it’s the supreme biological function . . . [*audience laughter*]

Santos: If I could trade my subjective experience with a jellyfish, I’m not sure I’d do it . . .

Paulson: I have a question for each of you, before we wrap up here.

If you could answer one question about consciousness, sort of the big question that preoccupies you, if there is one, what would it be? What's the question you most want answered? Dave, let me start with you.

Chalmers: I just want to understand it. So I guess the question is, what will a theory of consciousness look like? And maybe that's cheating because that question is so big, it's going to encompass answers to all the others. But, hey, I'm a philosopher . . . [*audience laughter*]

Kahneman: [*Pointing to Schiff*] I want to have his answers. I think that genuine progress is being made, that this type of work is science and they work with important facts; and they work on our intuitions and make for a better world. And we need more of this, I think.

Schiff: Well, I think my answer is the same as Danny's. But the question would also include something about a mechanistic account of how subjective experience arises in the brain. An answer to that would probably address, as David said, lots and lots of questions; and then we could do practical work with that model. Assuming that it's possible to construct a model, and I'd like to assume it is.

Paulson: Laurie?

Santos: I'm going to go with something else, which is that I would like us to have a full account of our intuitions about what is conscious, what's not conscious, and how all of these intuitions work at an information processing level. I believe this would make the world a better place because it would make our decisions better. Based on what we know about cognitive science and the methods that work pretty well, this is something that might be more likely to happen in our lifetime.

Kahneman: [*To Santos*] This is a feasible program: to understand our intuitions because they're pretty simple, though they're not internally coherent, which is why we can't use them as a basis for a science. But they do explain a lot about empathy. And so if we understood this about intuitions, a lot of social emotion would be understood at the same time.

Paulson: We could go on, but we are out of time. Thank you so much, David Chalmers, Daniel Kahneman, Laurie Santos, and Nicholas Schiff.

Commentary

The challenge of measuring consciousness: making sense of neuroimaging findings in the severely brain-injured

Nicholas D. Schiff

Many recent media reports have discussed the use of neuroimaging techniques to identify evidence of conscious awareness in patients with severe brain injuries. As covered in the panel discussion, despite major advances in neuroscience, recovery of consciousness after brain injury remains poorly understood. In this context, some preliminary remarks and vocabulary are worth presenting to the reader to aid understanding of this complex area. Perhaps the greatest challenge at present is the surprising degree of uncertainty of underlying brain function that may be present when confronted with patients at the bedside with very limited or even no overt signs of behavioral responsiveness. There is a range of behaviorally defined states of impaired consciousness observed after severe brain injuries beginning from coma (which is a state of unchanging behavior characterized by no evidence of response to the environment or self-generated behaviors), to a vegetative state (quite similar in showing no evidence of self or environmental awareness but distinguished by the cyclical appearance of an eyes-open or eyes-closed state), to a minimally conscious state (a condition with at least some evidence of awareness ranging from modest responses to sensory stimuli such as looking or head-turning to command following and attempts at communication), and up but not including patients in a locked-in state (verifiable bedside evidence of full consciousness with extremely limited motor control). Because motor disabilities may easily mask cognitive function or recovery in the setting of severe brain injuries, there are many patients whose level of consciousness, we cannot, at present, confidently assess. This state of current scientific knowledge has profound implications for tracking recovery after brain injury for applying treatments and formulating prognoses, as well as implications for measuring effects of new therapeutic inventions in the acute phase of treatment and recovery.

In the panel discussion an important distinction arises between patients identified as in the locked-in state and others who may retain such a high-level, normal cognitive function that may exist but cannot be verified. Recent studies have expanded our knowledge of the range of this uncertainty, highlighting the evolving picture of the natural history of the recovery from severe brain injury and the challenges of utilizing surrogate measurements with functional magnetic resonance imaging (fMRI) technologies or electrophysiological methods to determine evidence of awareness and higher cognitive functions (reviewed in Ref. 1). In evaluating and contextualizing much well-publicized research in this field it is important to recognize that identification of conscious acts indicate misdiagnosis of a patient as in an unconscious brain state rather than a new understanding of that brain state. Thus, patients in a vegetative state cannot imagine playing tennis and show such responses; only patients misdiagnosed as vegetative could. At present, it is not possible to estimate a number of such misidentifications. Although as many as 40% of patients may be misidentified as vegetative when fulfilling the diagnosis of minimally conscious state by expert examination, the vast majority of minimally conscious-state patients have not been able to carry out these tasks. The implication is that patients able to do high-level imaging in fMRI or electrophysiological experiments are likely to be closer to patients in a locked-in state with their capacities masked primarily by loss of motor control. Most importantly, such conditions should not be confused with brain death. Brain death can be accurately and unambiguously diagnosed and does not carry a prognosis. Rather brain death is death and is unrecoverable. As this area of medical science evolves, it will undoubtedly add important insights into the biology of consciousness and how science can make increasingly useful measurements of activity in brain that reflects the conscious mind. The interested reader may find more in the recent review by Laureys and Schiff.¹

Reference

1. Laureys, S. & N.D. Schiff. 2012. Coma and consciousness: paradigms (re)framed by neuroimaging. *Neuroimage* **61**: 478–491.