132

ABOUT DAVID CHALMERS, PhD

David Chalmers is a professor of philosophy and neural science at New York University. He received his PhD in philosophy and cognitive science from Indiana University in 1993 and has held positions at University of California, Santa Cruz, University of Arizona, and Australian National University. He helped to found the Association for the Scientific Study of Consciousness and co-directs the PhilPapers Foundation.

Chalmers has written extensively on a wide range of topics, including philosophy of mind, philosophy of language, metaphysics, epistemology, and metaphilosophy. He is a prolific author. His well-known works include *The Conscious Mind*, *The Character of Consciousness*, and *Constructing the World*. For more information about Chalmers and his current projects, visit http://consc.net/.

THINKING JUST HAPPENS

AN INTERVIEW WITH
DAVID CHALMERS, PhD

133

Stance: *How did you get your start in philosophy? Who or what influenced you, and how did that impact where you are today?*

David Chalmers: I started out in mathematics. My first degree was in math in Australia. I went on to do a graduate degree in math at Oxford, and then gradually got really interested in philosophy. I took one course in philosophy as an undergraduate at Adelaide in Australia. It was actually my worst course. I didn't do terribly well. But I guess it got me interested in the subject. That was my first year as an undergraduate. The rest of the time, I kept thinking about philosophy in the background, especially the problem of consciousness, which I gradually got slowly obsessed by. It always seemed to be that this was just the most interesting problem in the universe.

When I got to Oxford, I was meant to be thinking about math. But on my way to Oxford, I hitchhiked around Europe for about six months. I spent a lot of time standing by the side of the road thinking about philosophy. I got to Oxford and thought, "Alright, now I should get down to work." But, in fact, the whole time I was thinking about consciousness and the mind, and philosophy took over. It got to the point where I thought, "I should try and do this properly." Then I thought about doing philosophy at Oxford. In the end, I got in touch with Doug Hofstadter, whose book *Gödel, Escher, Bach* I'd read and really liked, as well as his book *The Mind's I*, which has a lot of issues about artificial intelligence and the philosophy of mind. I started writing back and forth with him, and he said, "You should come and work with me in Indiana." So I ended up moving to Indiana and working on philosophy and cognitive science there.

S: One of your books, The Conscious Mind, *was published in 1996 and another,* The Character of Consciousness, *was published in 2010. Has your view of consciousness changed through your career?*

C: Oh, there have been changes here and there over time. I started out as somebody very sympathetic to materialism. With my background in science, I thought that materialism was the most plausible view of the world, the default option. I thought consciousness ought to be something we can fit within science, that there ought to be a materialist explanation of it. I recognized that consciousness posed a big problem, but for a long time I thought we could find a materialist solution. Gradually, I came to the realization that given certain things I was committed to for systematic reasons, it couldn't be explained physically. So, earlier on, I ended up coming around from a more reductive, materialist view to a non-reductionist view of property dualism, even panpsychism.

I'd say I've changed in some smaller ways. My first book, *The Conscious Mind*, was very sympathetic to epiphenomenalism: the view that consciousness doesn't play a causal role. I'm now somewhat less sympathetic to that view. I'm especially interested in views in which consciousness does play a role—interactionist dualism, maybe panpsychism. At the same time, I've also become very interested in all kinds of views, whatever view it might take to solve the mind-body problem. I've even had a bit of resurgence of interest in a view on the

extreme opposite of these views, which is illusionism: the view that consciousness is basically an illusion. Dan Dennett might have a version of this view. Recently, philosophers like Keith Frankish have begun to explore it.

I'm actually writing a paper at the moment on what I call "the meta-problem of consciousness," which is the problem of explaining why we think there's a problem of consciousness. In the past, I've distinguished the easy problems—the problems of behavior—from the hard problem—the problem of explaining experience. But here's one bit of behavior which is very closely tied to the hard problem: people like me go around saying things like, "Hey, there's a hard problem. Consciousness is hard to explain. I can't see how it could be physical." That's a bit of behavior that we could try to explain in physical or functional terms. I'm getting really interested right now in the problem of how we might actually explain that in physical or functional terms. I talked about it a bit in my first book, *The Conscious Mind*, but I think there's more to say.

Once you go that far, then some people will be very inclined to say, "Ah, what you've really done is explain the illusion of consciousness, given a physical explanation of why you think consciousness is non-physical." For various reasons, I don't think that's the correct conclusion to draw. But I do think that coming up with a really good explanation of why it is that we say and judge these things about consciousness is sure to tell us something about the basis of consciousness itself.

S: What are the principal aspects of consciousness that you want to keep? What are some core properties that you see as most important?

C: Oh, interesting. I think the most important one is that consciousness is a state that it's something like to be in. My colleague Tom Nagel, 40-plus years ago now, in his article "What is it Like to be a Bat," said you're conscious if there's something it's like to be you. A mental state is conscious if there's something it's like to be in that state. I take that to be the defining feature of consciousness.

Philosophers sometimes talk about phenomenal consciousness to distinguish it from other things that people call consciousness. You can put it in different ways by saying it's subjective or it's a state of experience, but that's really the core thing. At the same time there are many other features that consciousness obviously

has. It's multi-faceted. There are different kinds of consciousness, different sensory modalities. It appears to present the world as representational; it seems when I'm conscious, I'm conscious of a world around me. That seems like a very important feature of consciousness. It's integrated in various ways, and it seems to be unified. But it's also differentiated in various ways. Up to a certain point, I don't want to be too rigid in laying down exactly what the features are of consciousness. We could turn out to be wrong about some of them. Is it completely representational? Is it always? Well, I don't know. Is it unified? I'm inclined to think so, but maybe someone could convince me otherwise. I guess the really important thing is just that there's something it's like to be in that state, and everything else we can argue about.

S: On the same path here, are there different levels of consciousness throughout different animal species, like elephants or dolphins? Or is there any fluctuation in how they experience consciousness, or how they experience the world, or how they are "being," I guess, as you put it?

C: We just had a conference three weeks ago here at NYU on animal consciousness, organized by the Center for Mind, Brain, and Consciousness that I direct along with my colleague Ned Block. It was quite an eye-opener, this conference. People used to argue about whether certain mammals are conscious. Maybe they thought that other primates were conscious: dogs and cats, maybe. Mice? It starts to be doubtful. These days people are much more open to ascribing all kinds of consciousness to animals. Pretty well every mammal; it's regarded as pretty well obvious. There used to be people arguing that fish don't feel pain, but I think that's gradually becoming a minority position. The debate has now moved to insects, actually. People are now arguing about whether ants or bees are conscious. The tide seems to be gradually going in that direction. I think a lot of different animals are very plausibly conscious.

It's predictable that I might think this. I've got some sympathy with panpsychism, the view that everything is conscious. If everything is conscious, then animals are conscious. In general, most people seem open to ascribing consciousness to a lot of animals. But that's not to say they all have the same kind of consciousness, so maybe that's where your points about levels of consciousness come in.

Maybe all these animals have some kind of sensory consciousness. Most of them have some kind of vision. Maybe they have some kind of visual consciousness, but even within visual consciousness there might some serious differences. Some of them might have sophisticated color vision. Some might lack it. Some might be just sensitive to brightness.

I think human consciousness involves much more than just sensory consciousness. We have cognitive consciousness when we think. There's a phenomenology of thinking. We also have some kind of consciousness of ourselves, reflective consciousness, reflecting on our thinking. You can call those, if you like, higher or more complex levels of consciousness. I'm pretty doubtful that ants or bees have anything like that. Maybe they have simple visual consciousness. Maybe they have simple feelings of pain. I'm pretty sure that they don't have reflective consciousness, thinking about their thinking. Maybe they don't even have cognitive consciousness, a phenomenology of thinking. But they may well have sensory consciousness. Then it's a very interesting question, what consciousness different non-human animals might have. Do non-human primates, for example, have some kind of consciousness of thinking? Smarter primates are very sophisticated. But I think we're still at the point of really just charting out these things. We certainly haven't yet figured out what it's like to be a bat, what it's like to have a sensory modality totally alien from human sensory modalities.

S: Are there ethical implications if panpsychism, or this idea of consciousness spread throughout the world, is true?

C: In general, the more you extend consciousness to more creatures, the more you're including them within some circle of ethical concern. If fish are not conscious, that's very convenient to the fisheries because then they don't need to worry about causing them pain. Presumably, then, there's no real issue about killing and eating them because a lot of people think consciousness is what makes a creature come to warrant certain kinds of moral concern. Insofar as we acknowledge that fish are conscious, then we could be thinking about things like: Are they suffering? Is it appropriate to cause them pain? Is it okay to kill them? Likewise for ants and so on.

By the time you're done extending this to everything, which is the panpsychism view—that everything is conscious—you might say, "Well, extending the circle of ethical concern to every object in the universe

becomes a little bit ridiculous." Think about all the atoms I'm messing around with when I drink a glass of water. Am I abusing those water molecules? It starts to become a bit strange sounding if you should extend your ethical concerns to everything in the universe.

As I've become more sympathetic with panpsychism, I've started thinking there's more than just consciousness that makes something an object of moral concern; it's some forms of consciousness. Some people think that it's pain or suffering that matters especially, but maybe it's something more than mere sensory consciousness. Maybe it's some type of reflective consciousness or cognitive consciousness that gives a being a greater level of moral status.

I used to say that I shouldn't eat anything that was conscious, and therefore I should become a vegetarian. I'm not vegetarian, but I thought that maybe I should become vegetarian. Then the more and more you become sympathetic with panpsychism, then the more and more things you're not going to be able to eat, until you're left with being able to eat nothing at all. But the consequence of that to me was that it's not just consciousness that matters, but that it's certain kinds of consciousness that matter. But I've certainly not yet thought it all the way through.

S: In "Extended Cognition and Extended Consciousness," you describe consciousness as a subjective experience that is "one step away" from moving the body. Could you say more about what this step is and what the implications of its existence are?

C: The thought was that there is a distinction between conscious processes and unconscious processes. The traditional thing to say is that conscious processes are the ones that are reportable. The conscious processes are ones that you can talk about. For example: "I'm experiencing red," "I'm feeling pain." The unconscious stuff might be in the brain somewhere, but it's much less accessible. I was picking up on that idea and generalizing it a bit, because we don't want to make reportability define what consciousness is. We want to say there are animals which can be conscious without reporting, but then what will be the conscious processes in them, the processes which are somehow directly available to the creature in certain ways for controlling their action and for controlling other processes and so on? So phenomenal consciousness will still correspond

to what's accessible. At the level of processing, it corresponds to what's accessible or available. That was something I argued for, maybe 20 years ago, in an article called "Availability: The Cognitive Basis of Experience." But it also came up in "The Extended Mind."

I'm sympathetic with the idea that some mental states can extend beyond the head, like our beliefs, some memories. I'm inclined to say that many of my memories are now stored in my iPhone. Andy Clark and I argued that many mental states can extend. But it seems to me that it's mostly not conscious states, but more dispositional states like belief, which extend beyond our consciousness. Then the question arises: can you get consciousness to extend beyond the head? Maybe you can in science fiction cases by wearing a module or a belt or something. But it's harder to see how it's going to go in regular cases involving perception and action. I have never seen a convincing case where my consciousness extends into my iPhone. The question is: why is that? Well, maybe it's because the iPhone is too far removed from the locus of control. You said one step away; maybe the iPhone is three or four steps away from that, so it's only indirectly available for control. The consciousness corresponds to the stuff that's one step away; it's poised for control. If consciousness has to go with what's poised for control, then it's pretty plausible that the processes which are poised directly for controlling behavior are in my head. That's why there's no extended consciousness, because the stuff which is outside of the head is too many steps away to count as directly available for control.

S: Throughout your career, you've been writing about possible worlds and two-dimensional semantics. How would you simply define and describe your work to someone who hasn't quite encountered those terms before?

C: I think of it as a project which is broadly in the spirit of trying to make sense of Frege's distinction between sense and reference. So you've got "Hesperus," the word for the Evening Star, and "Phosphorus," the word for the morning star. Frege wanted to say they've got the same referent—the planet Venus—because it's Venus both times. There's something in the meaning of those words that's the same: the referent. But there's something in the meaning of those words that's different, and that's what Frege calls the sense. One sense went along with "the morning star," and

one sense went along with "the evening star." Two-dimensional semantics is trying to split meaning into two components in ways similar to the way that Frege did it.

But many people think that the Fregean picture is made problematic by the work of Saul Kripke in *Naming and Necessity* around 1970. He argued that names are rigid designators, that the meaning of all these terms is just the planet Venus. He said various descriptivist intuitions, which are broadly Fregean intuitions, were wrong for various reasons.

Some people, like Carnap, had wanted to make sense of Fregean sense in terms of what a word picks out in different possible worlds. An intention, for Carnap, in the case of "Evening Star" would pick out in any given world the evening star in that world and "Morning Star" would pick out the morning star in that world. In our world, they might be the same—have the same extension—but in other worlds "Morning Star" and "Evening Star" might pick out something else and have a different intention. This was Carnap's way of making sense of Frege's intention and extension.

Then Kripke comes along and says, "Actually, if you look at 'Hesperus' and 'Phosphorus,' they're rigid designators." They pick out the same object in all possible worlds. If you want to look at the intention for "Hesperus" or "Phosphorus," you have to look at them as rigid designators. "Hesperus" is Venus in every world. "Phosphorus" is Venus in every world. More generally, he argued for anti-Fregean, anti-descriptivist conclusions, which led many people to think, "Oh, really, all there is here is reference."

So, two-dimensional semantics is partly in reaction to Kripke. It's a way of splitting meaning into two components, one of which behaves roughly the way that Frege and Carnap thought meaning behaves—the way that's like sense. The other dimension corresponds to the way that Kripke thought meaning behaves. So, technically with possible worlds, both aspects of meaning are represented as intentions, as functions from possible worlds to extensions, but you've got two intentions now. You've got one that behaves the way Kripke thinks, where "Hesperus" picks out Venus in all worlds. That's the secondary intention. But you've got another one that behaves more like the way that Carnap said meaning behaves, where the intention for "Hesperus" will pick out the Evening Star in every

world and the intention for "Phosphorus" will pick out the Morning Star in every world. So we'll say that "Hesperus" and "Phosphorus" have the same secondary intention—that's one dimension of meaning—but a different primary intention—another dimension of meaning. So Kripke's modal argument against descriptivism—that names are rigid designators—gets accommodated by the secondary dimension of meaning. But all the Fregean intuitions about the cognitive differences between "Hesperus" and "Phosphorus," or "Superman" and "Clark Kent," get accommodated by a difference in the primary intention for those things.

There's a lot more to say about how this plays out in specific cases and deals with all of Kripke's challenges, but the broad idea is that we've got two dimensions of possible worlds, which for me come down to a kind of epistemic possibility and a kind of metaphysical possibility. The epistemic one corresponds to the primary intention and the way words behave in respect to those go with Frege's sense. The secondary one, which is over metaphysical possibilities, behaves Kripke's way, much more like a notion of meaning tied to referent. It's a way to try and have your Frege and your Kripke, too.

S: *In "The Two-Dimensional Argument Against Materialism," you described phenomenological concepts like consciousness that are metaphysically necessary. Are there other things that are metaphysically necessary besides consciousness?*

C: I don't think I used that exact thesis. I would say that consciousness is contingent: it exists in some worlds, but not in others. I did say some things about concepts of consciousness being special, maybe being super rigid, and epistemically rigid.

Kripke says that names are rigid designators. By my lights, they would be called metaphysically rigid designators. If we think about what is "Hesperus" in different metaphysically possible worlds, Kripke would say that it's always the planet Venus. The secondary intention always picks out the same object. But if we do it across epistemically possible worlds—say we're actually living in a world where Jupiter is visible in the evening and Mars is visible in the morning—then, relative to that epistemically possible world, if things are actually that way, we'd say "Hesperus" picks out Jupiter and "Phosphorus" picks out Mars. That is to say, "Hesperus" and "Phosphorus" are metaphysically

rigid. Their secondary intention pick out the same object in every world. But they're not epistemically rigid because we can still think of epistemic possibilities where "Hesperus" is not "Phosphorus." Kripke said it's metaphysically necessary that "Hesperus" is "Phosphorus." But it isn't epistemically necessary. It's not *a priori* that "Hesperus" is "Phosphorus" or that "Clark Kent" is "Superman." So we can say that these names are metaphysically rigid, but not epistemically rigid.

Now there are a few special expressions, I think, which are not just metaphysically rigid, but epistemically rigid, too. I think one plausible example for this is number terms, like "the number nine." Maybe that picks out the number nine in all metaphysically possible worlds and the number nine in all epistemically possible worlds, too. You don't get this Kripke thing where you don't really know *a priori* what it refers to. You know *a priori* what it refers to; it's epistemically rigid. So maybe number terms, like "nine" or "zero," maybe some special property terms are dually rigid.

I think "consciousness" is like that. It's epistemically rigid. When we think about "consciousness," it's not like "water" or "Hesperus," where it's picking out something out there in the external world but we don't really know what. My view is that we grasp the essence of consciousness simply by possessing the concept. In every possible world, even in every epistemically possible world, "consciousness" picks out consciousness. That's to say it's not just metaphysically rigid, but it's epistemically rigid. And that makes the concept of consciousness somewhat special.

I call this kind of epistemic and metaphysical rigidity "super rigidity." That kind of phenomenon comes close to what Russell was talking about under the heading of acquaintance. There are some special things we can be acquainted with. He thought we could be acquainted with sense data. He thought we could be acquainted with certain universals, and he thought we could be acquainted with the self. I'm inclined to think that the cases where we have that similar kind of acquaintance correspond very nicely with the cases where you have this epistemic rigidity. I think of this epistemic rigidity as a way of trying to make sense of this Russell–style notion of acquaintance but in the more contemporary framework of possible worlds.

S: Just to switch gears just a little bit: in your own words, what does it mean for something to be scrutable? How does this relate to our reality and other branches of philosophy like ethics?

C: This was the notion that was the centerpiece of my book, *Constructing the World*, which came out about five years ago. Scrutability is the idea that once you're given a full enough description of the world, you can figure out all the truths about the world. One way to introduce this is by thinking about Laplace's Demon. Laplace said, "Tell me all the laws of physics and tell me all the positions of all the particles at the beginning of the universe, then I will be able to predict the entire future of the universe and I'll be able to know everything, not just about atoms but about chemicals, about organisms, about society, and so on." Laplace held that everything about the world was scrutable from the facts about the position of the particles. Scrutable here means basically given the knowledge of A, you can figure out B. The key notion I talk about is developed in terms of the *a priori*. Given knowledge about A and *a priori* reasoning, you can figure out B. It's a kind of deducibility, if you like.

Now Laplace's thesis is often taken to be too strong in that there's various things that are hard to know. If physics isn't determinist, then maybe you can't even know about the particles in the future. Many people have argued that knowing anything about physics doesn't tell you about the conscious experiences. John Perry's got this point that knowing objectively the full state of the world doesn't tell you which person is me. But I suggest that those are all gaps we can close by building more things into the basis. If we have a broader basis, which is all the physics and all the facts about physics and all the facts about consciousness and maybe an indexical fact about where I am—it's a big truth I call PQTI: physics, qualia, T is "that's all" (or "this is everything"), and indexicals—from there, you can deduce everything. Everything would be scrutable from there. That's what I try to argue for in *Constructing the World*.

I also try to argue that you can in this way vindicate some of the projects of the logical empiricists like Rudolf Carnap, who in *The Logical Structure of the World* tried to argue that you could get all truths about the world in a very small vocabulary. He actually thought he could just do it with logic and a little bit of a relation called phenomenal similarity. I think, like Laplace, Carnap was a bit too ambitious with his minimal vocabulary,

but then I try to see how little can we expand the vocabulary and still get something from which we can deduce all the truths. I thought, maybe from PQTI, we can do it. Later in the book I try to reduce that down and define a lot of the physical terms in terms of mathematics and causation, and not much else. By the end of it we do have a very simple set of fundamental concepts in terms of which we ought to be able to describe the world fully, so that from there one could deduce everything about the world.

I think it's an interesting question, how ethics fits in there. It partly depends on whether you're a moral realist, because I do all this for truths about the world. If you do believe that moral claims are true or false, I think it's pretty plausible they're also *a priori* deducible from non-ethical claims. Given a non-ethical description of a scenario and all of the experiences and so on—I don't know, something awful like Harman's case of torturing a cat—it seems pretty clear that you can figure out, "Okay, that's bad." So maybe the ethical claims about the world will be scrutable from the underlying claims.

Lately, I've actually been thinking about whether there might be stronger forms of scrutability where it's not merely *a priori* deducibility but analytic deducibility. Arguably that's missing in the ethical case. One question is: what is the connection between this notion of scrutability and existing notions, like supervenience or metaphysical grounding? Scrutability goes along with supervenience, which is a kind of necessitation under certain circumstances, but is it enough for reduction? The mathematical truths are scrutable. They're *a priori*. They're scrutable for everything. Does that mean they're reducible to the physical truths? Maybe not. The ethical truths, they're scrutable for the physical truth, but does this mean they're reducible? Maybe not. Maybe for that stronger connection you need something like analytic scrutability or analytic entailment. And maybe that missing in the ethical case. So, that's a thesis I've been pursuing more recently.

In general, a lot of my broad life project is to connect certain metaphysical notions, like metaphysical necessity or grounding, with certain epistemological notions, like *a priority* or closely related notions like analyticity, to draw connections between them that are stronger than what other people think are possible. I think scrutability can actually give you an epistemological lens of some of those metaphysical questions about the ultimate structure of reality.

S: Along the lines of more two-dimensional semantics and conceivable worlds, what kind of things are inconceivable, do you think? Or do you think with the PQTI method that we can deduct most things and know most things?

C: I think that most obvious things are very inconceivable. Consider things which are ruled out *a priori*. Can I conceive that two plus two is five? I don't think that I can conceive that two plus two is five. Can I conceive of a round square?

Well, I don't really think so. Roy Sorensen, a philosopher, thinks he can draw a round square. But you can only draw a round square from edge on. So it looks just like a line. That's probably cheating when it comes to conceiving of a round square. So the most obvious things which are inconceivable are things ruled out *a priori*.

Some of my colleagues think that zombies, physical duplicates of us without consciousness, are inconceivable, but I find them pretty well conceivable. They would say a physical duplicate of me that was not alive—that was functioning just like me, metabolizing and adapting, reproducing, but not alive—I can't really conceive of that. That would be a case where, when something like life is scrutable from all the underlying facts, it ought to be difficult beyond a certain point to conceive that all the underlying facts without the high level thing, like life. That would again be an example of something by my view that is ruled out *a priori*.

I think what is interesting is there are some other things that are not ruled out *a priori*, which are possible in principle but nonetheless are inconceivable. Probably for limited beings like us, there are such things. For example, what it's like to be a bat. If we don't have the bat's sonar system in our head, we can never form a really detailed conception of what it's like to be a bat. So that's inconceivable for us, but maybe it's not inconceivable in principle. Some being could conceive of it. Maybe a bat, or a souped-up bat, or a future version of humans where we have a bat sonar module plugged into us. The harder question is: "is there something which is possible but inconceivable in principle, some properties of the world that no being could possibly conceive of?" I'm enough of a rationalist that I would like to think that's impossible, but I have to admit that I have no argument. I'd like to think that

everything about the world is at least possibly intelligible to some being, but I can't say I've got any good argument for that. I think that's still an open question.

S: Switching gears, what changes have occurred in your overall philosophy since you started publishing? What ideas from your early career would you work on, improve on, or disagree with completely? And what do you think is your trajectory of your life's work?

C: I've gradually gotten more and more interested in more areas of philosophy. I started out just interested in the mind-body problem. I got into philosophy because I was super interested in the problem of consciousness. I thought it was the hardest problem in the world, the most interesting thing to figure out. At the beginning, it wasn't philosophy in general but that problem in particular. But pretty quickly I figured out that, to get a grip on some of these metaphysical problems of consciousness, I've got to think about metaphysics. You start to get at metaphysics, reduction, grounding, supervenience, properties, whatever. Then, to get a grip on that, you have to start thinking about the philosophy of language—the language we use in thinking about the stuff, these Frege and Kripke points, and so on. So I started thinking about that. Then I started thinking about epistemology, and then you end up thinking about the philosophy of science and even connections to some issues in ethics, especially metaethics. I started thinking about philosophical reasoning in general. For me, the experience has been one of opening out and opening out into more areas, although an interest in consciousness has remained at the core of what drives me in the end. So, almost all of philosophy has become fascinating. I haven't written a paper in political philosophy yet or in aesthetics. Maybe one of these days. We'll see.

Any idea I wrote about when I was younger that I disagree with now? That's interesting . . . Philosophy for me also has gotten more and more systematic. Everything connects to everything else. I started out thinking I was a rationalist or empiricist—some kind of broad rationalism, that the universe is fundamentally intelligible. That's part of what drives my work to get deep links between epistemology and metaphysics. I think of myself as a philosophical optimist. These problems are ultimately solvable even though we might

not have solved them yet.

I ought to have some idea I wrote about when I was younger that I disagree with now. There's a lot of specific things I would disagree with. There are few claims here and there that were just mistakes that I have corrected over the years. But I think in terms of broad, general views, it's not like I've done a U-turn. It's not like the late Wittgenstein, who totally rejected the early Wittgenstein. Maybe I still have time to become the late Chalmers who rejects and repudiates the early Chalmers. As I said, I'm getting interested in illusionist views, where consciousness is an illusion. Maybe there's time for me, in five years time . . . I did think about, just for fun, publishing a book called *The Inescapable Illusion of Consciousness*, repudiating everything I ever said about consciousness in the past, but I haven't quite come to the point yet where I can believe it. Maybe it'll happen.

S: Where do you think philosophy as a discipline is headed? Will a new field—or a combination of fields—emerge, similar to your conjunction of philosophy of language and philosophy of mind?

C: It's hard to say. New fields of philosophy do emerge from time to time. What happens much more often is that a new subfield emerges and new topics emerge to think about. One thing that's gotten extremely big in recent years is the philosophy of technology. People are thinking about computers. For me, technology generally, and computers specifically, have led to a huge amount of philosophical enrichment. It's hard to think about the philosophy of mind these days without thinking to some extent about artificial minds and artificial intelligence. Likewise for metaphysics. Thinking about reality, thinking about metaphysics of the world, you also you want to think about artificial reality, virtual reality. People are now spending more and more time hanging out in virtual worlds. Virtual reality devices are just starting to go pretty wide. They are going to play a very major role in coming years. A lot of my work lately has been thinking about virtual reality and virtual worlds. More generally, thinking about technology poses any number of new philosophical questions: the ethics of technology, how we live with technology, the role of computers in our society, and how to handle that in a fair and just way.

I also think there has been a movement over the last ten years or so for philosophy to get more deeply connected

to social issues, whether it's issues about race or gender and sexuality, issues about technology, issues about social epistemology, issues about the metaphysics of the social world. I have noticed a really big trend in that direction. Is that a wholly new area of philosophy? I'm not sure. But I think it's one of the areas that is the richest and most exciting right now, using our philosophical methods to think about the social world and the technological world.

S: *How can philosophers successfully engage the public on more sophisticated topics in philosophy?*

C: I think some philosophy is technical. There is no way I'm gonna be able to do two-dimensional semantics, as is, in a newspaper article. I could give some examples and give a very simplified version of it, but some part of the interest of some ideas is essentially technical. But that said, even where there are technical ideas, if they are important ideas, there ought to be a way of putting some of the central points and conveying some of the key ideas which is not technical. Over the years, I think in my work on consciousness, I've pursued it at both levels. Some of the time it's detailed. It involves supervenience; it involves two-dimensional semantics; it involves really finicky detail, but I think that's important. At the same time, I care a lot about communicating to a broader group, whether it's academics in other fields, such as scientists thinking about consciousness, or just intelligent people from any walk of life who are interested in the mind. Using more general language: here's the hard problem of consciousness compared to the easy problems; here are the kinds of ideas you might have available. You might need new fundamental properties like consciousness. Maybe they'll need to be everywhere, maybe not. These are the big broad issues which are driving the technical discussion too, except the technical discussion often takes place at a much finer level of detail.

Actually, I wrote an article fifteen years ago on *The Matrix*, on whether we could be living in *The Matrix* and what we can know. I tried to argue that even if we are in a matrix, the world around us is still real. That article was presented in a very simplified way. Maybe some people think, "Well, he's not being serious, he's slumming it by talking about *The Matrix*." But I think I was presenting an original article for a really important and serious philosophical view. I recently wrote a more detailed piece, trying to fill out the details for this for a philosophy journal article. I called it "Structuralism as

a Response to Skepticism." It doesn't hurt to add a few more details in it and a bit more academic rigor about it. But I don't want to say that makes it less serious. Even the version which is cast in terms of *The Matrix* can convey fifty percent of the philosophical content. It's just a matter of trying harder to be clear and not using jargon in talking to the public. It helps to give it a hook. Some kind of story or something of current interest helps get people engaged. But my experience is that people in the public are just really interested in philosophy. Now maybe it's the selection effect: people I meet are the ones who are interested in philosophy. But there's a hunger out there for philosophy, and anyone who's willing to talk to them in a serious way that is still clear and accessible. . . . I think the public will be interested in that.

S: *What is the most useful advice you would give to anyone in pursuit of a degree in philosophy?*

C: I think you have to be really passionate about philosophy. It's not easy to get a job in philosophy, and inevitably a lot of people who start graduate degrees are going to not end up with permanent positions in philosophy. So I would say, at the very least, be passionate enough about the value of philosophy so that if you don't end up getting a permanent academic job in philosophy, you'll still be happy you've done the graduate work. Do it, by all means. If you're seeing this as a meal ticket, there's probably a bit of a problem.

For all the people who are passionate about philosophy and want to think about it, well then it's very worthwhile for them to go to a graduate school. If you want to think about all this full time, I'd say try to develop your skills. It's very important to develop skills, tools, and knowledge of literature, and all those professional things. But I also think it's important not to lose sight of the things that made you passionate about philosophy in the first place. If you've got a burning desire to pursue a certain idea, even though it's not the fashionable thing or if it involves going a certain direction, don't be afraid to follow your own original idea that made you passionate in the first place. I think we always need new ideas in philosophy. The field right now is ready for new ideas. I'd love to see smart people going into graduate school and developing all those tools and becoming professionally aware of the literature. Those things are important. But still using the skills they develop to follow original ideas and new ideas and bold ideas, that's how philosophy gets transformed. Every now and then

IT'S IMPORTANT NOT TO LOSE SIGHT OF THE THINGS THAT MADE YOU PASSIONATE . . . IN THE FIRST PLACE.

a Frege comes along, or a Wittgenstein comes along, and follows their own views in a new direction; I'd like to see some people going on to graduate school in philosophy doing that.

S: So last question: do you ever hitchhike anymore to have time to think?

C: I've not hitchhiked in some time, maybe once a few years ago when I ran out of gas. When I moved to the U.S. in the late eighties, hitchhiking was not such a big thing. People didn't think it was especially advisable. But it is hard finding time to think. As an academic, your time can easily be eaten up by the next thing you need to do, the class you need to teach, the committee you want to serve on, the articles you're about to review, endless emails you have to answer . . . so, it's not easy. Those things are all fine. But I think it is important to carve time to think and to work. Living in New York, which is where I am now, the city is great, but there are endless distractions. So one thing that's made a difference to me is having a place out in the country, outside of the city, where I go up to and try to use that mostly for thinking and writing and chilling to some extent. It's a place to get away from all that bustle of all the constant demands of day-to-day academic life and to withdraw and think about new ideas, and especially to write them down.

Actually, to be honest, the thinking is not the hard part for me. Thinking just happens during the day, in the shower, walking to work, whatever. The thinking is always going on. It's the writing that's really the bother. Finding time to not just think, but actually get things fully formed and written down, that's the hard part. That's the part where it's nice to have time to do it. We're just coming up on the end of the semester now. Once the semester ends and we get through the holidays, I'm looking forward to going up to that place up in Hudson Valley, and having time to think and write and to get stuff properly on paper.