
*Conversations in the Cognitive
Neurosciences*

edited by Michael S. Gazzaniga

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

10

Qualia

Daniel C. Dennett



Daniel Dennett was educated at Harvard and Oxford, receiving his D. Phil. in 1965. After six years at University of California, Irvine, he moved to Tufts, where he is Distinguished Professor of Arts and Sciences and Director of the Center for Cognitive Studies. He is the author of articles on many issues in artificial intelligence, psychology, and cognitive ethology, as well as in philosophy. His books are Content and Consciousness (1969), Brainstorms (1978), The Mind's I (with Douglas Hofstadter, 1981), Elbow Room (1984), The Intentional Stance (1987), and Consciousness Explained (1991). His latest book is Darwin's Dangerous Idea (1995).

MG: You are known both as a philosopher and as a cognitive scientist. How do think of yourself? Where does the one role stop and the other start?

DD: I consider myself a philosopher. Before the twentieth century philosophers often became quite embroiled in the science of their day (with mixed results, of course!), so my involvement with the details of cognitive science is not such an anomaly as it may appear when it is contrasted with the more recent stereotype of the philosopher who just sits in his armchair and claims to figure it all out from first principles.

Philosophy of science is one of the strongest—I think the strongest—of the subdisciplines in philosophy these days, and there are philosophers of physics who are quite at home in the lab or the farthest reaches of theory, philosophers of biology whose contributions mingle fruitfully with those of the more theoretically minded evolutionists, and so forth. I am trying to do the same thing in cognitive science. My goals and projects differ in two ways from those of some other philosophers working this territory.

First, unlike some philosophers of cognitive science, I do not view my role as solely what we might call “meta-criticism”—analyzing and criticizing the theories, arguments, and concepts of the scientists. On the contrary, I aspire to create, defend, and confirm (or disconfirm) theories that are directly about the phenomena, not about theories about the phenomena. The philosophers’ meta-criticisms are often important clarifiers and exposers of confusion, and as such are—or should be—unignorable contributions, but I myself would also like to make more direct contributions to theory.

Second, and following from this, I don’t consider cognitive science to be simply a mine from which philosophers of mind can extract valuable support for their purely philosophical theories. It is that, of course, and the insights gleaned from cognitive science have transformed—if not quite killed—traditional philosophy of mind. But what philosophers of mind sometimes fail to appreciate is that the scientists are just as susceptible to conceptual confusions as the “layman” and hence the fruits of their research cannot be taken neat and used as a stick to beat sense into the benighted layman. There are at least as many closet Cartesians and uncritical believers in “qualia” among the scientists as among the uninitiated, for instance, and these scientists have something to learn from philosophy (whether they like it or not!).

I don’t have a lab or do experiments, but I do devote a lot of effort to proposing experiments (or perhaps I should say “provoking” experiments) and to redesigning and criticizing experiments. And I have discovered, of course, that there is no substitute for direct experience in the lab. Many times I have thought I understood a series of experiments from reading the literature on them, only to uncover a fairly major misapprehension on my part when I actually witnessed the paradigm, or became an informal subject. Live and learn. That’s why, although I am a philosopher, not an experimental scientist, I can’t do my work well without poking my nose in the labs. Besides, it’s much more interesting than just reading philosophy journals.

MG: So in a sense the philosopher’s role is to prevent thought disorders among scientists. Likewise a simple empirical fact can raise havoc with a philosopher’s theory of mind, requiring that the philosopher keep abreast of recent discoveries. Before going further, can we get out on the table what you mean by qualia?

DD: I thought you’d never ask. Qualia are the souls of experiences. Now, do you believe that each human experience has its own special and inviolable soul?

MG: What are you getting at? What on earth does that even mean?

DD: That’s just my wake-up call for people who think they know what qualia are. It’s frustrating to learn that in spite of my strenuous efforts, people keep using the term “qualia” as if it were innocent. Consider a parallel: According to Descartes (and many churches), the difference between us and animals is that animals have no souls. Now when Darwin showed that we are a species of hominid, did he show that there really aren’t any people after all—just animals? If Darwin is saying we’re just animals, he must be denying we have

souls! So he must be saying that people aren't really people after all!

That's silly, but it isn't as if we didn't sometimes talk that way:

"You're behaving like an animal!"

"But I *am* an animal!"

or:

"They treated us as if we were animals."

In spite of tradition, the very real and important differences between people and (other) animals are not well-described in terms of the presence or absence of souls fastened to their brains. At least I would hope most of your readers would agree with me about that. Similarly, the differences between some mental processes and others are not well-described in terms of the presence or absence of qualia—for what are they? Not only is there is no agreed-upon definition among philosophers; controversies rage. Until they get settled, outsiders would be wise to avert their gaze, and use some other term or terms—some genuinely neutral terms—to talk about properties of subjective experience.

In fact the term "qualia"—which is, after all, a term of philosophical jargon, not anything established in either common parlance or science—has always had a variety of extremely dubious connotations among philosophers. Denying there are qualia is more like denying there are souls than like denying that people are much smarter than animals. If that makes "qualia" sound like a term one would be wise to avoid, good!

To put it bluntly, nobody outside of philosophy should take a stand on the reality of qualia under the assumption that they know what they're saying. You might as well express your conviction that trees are alive by saying they are infused with *élan vital*. So when Francis Crick, for instance,

says that he believes in qualia, or when Gerald Edelman contrasts his view with mine because his view, unlike mine, allows for qualia, these pronouncements should be taken with more than a grain of salt. I'd be very surprised if either Crick or Edelman—to take two egregious examples—believes in what the philosophical fans of qualia believe in. If they do, they have a major task ahead of them: sorting out and justifying their claims against a mountain of objections they've never even considered. I would think they'd be wise to sidestep the mess.

I fear I'm losing the battle over the term "qualia," however. It seems to be becoming the standard term, a presumably theory-neutral way of referring to whatever tastes and smells and subjective colors and pains are. If that's how it goes, I'll have to go along with the gang, but that will just make it harder to sort out the issues, since it means that all the controversies will have to be aired every time anybody wants to ensure that others know what is being asserted or denied. Too bad. Don't say I didn't warn you.

MG: Well, OK. These things happen. "Qualia" is doomed to mean the feeling about the specialized perceptual and cognitive capacities we humans enjoy. Put directly, should we not distinguish between the task of characterizing the cognitive operations of the human mind and the (here we go) qualia we have about them?

DD: Certainly we should divide and conquer. So we should distinguish between the task of characterizing *some* of the cognitive operations of the human mind, and the rest (which we conveniently set aside till later); but if we call the latter "qualia" and think that they are somehow altogether different from the "cognitive operations" we are studying now, we prejudge a major question.

Take experienced color, every philosopher's favorite example of a quale. Suppose what interests you as a cognitive scientist are the differences in people's responses to particular colors (Munsell color chips will do for standard stimuli, at least for this imaginary example). But instead of looking at such familiar measures of difference as size of JNDs, or latency of naming, or choice of color words (where does each subject's "pure red" lie on the spectrum, etc.), or galvanic skin response, or some ERP difference, suppose you looked at variations in such hard-to-measure factors as differences in evoked memories, attitude, mood, cooperativity, boredom, appetite, willingness to engage in theological discussion . . . you name it. Until you've *exhausted* all these imponderable effects, you haven't covered all the "cognitive" or "disposition-affecting" factors in subjective color experience, so there will be features of color experience, features of what it is like for each individual, that you are leaving out of your investigation. Obviously. But if you then call these unexamined residues "qualia" and declare (or just assume) that these leftovers are somehow beyond the reach of cognitive science, not just now but forever, you are committing a sort of fallacy of subtraction. There need be nothing remarkable about the leftovers beyond their being leftovers (so far). When some qualia freak steps up and says, "Well, you've got a nifty account of the cognitive side of color vision, but you still have a mystery: the ineffable what-it-is-likeness of color qualia," you needn't concur; you are entitled to demand specifics.

To cut to the chase, I once got Tom Nagel in discussion to admit that given what he meant by "qualia," there could be two identical twins whose scores on every test of color discrimination, color preference, color memory, effects of color on mood, etc., came out the same, and there would *still* be a

wide-open question of whether the twins had the same color qualia when they confronted a particular Munsell chip! (By Nagel's lights, neither twin would have any grounds for supposing that now he knew that he and his twin brother had the same color qualia.) Nagel's position is an available metaphysical position, I guess, but I hope it is obvious that it doesn't derive any plausibility from anything we have discovered about the nature of color experience, and hence no cognitive neuroscientist needs to be shackled by any such doctrine of qualia.

By the way, this should make it clear why I said qualia were the souls of experiences. Nagel's position is parallel to that of the vitalists of yore who, after being shown all the details of metabolism, biochemistry, etc., still held out that Life was not being accounted for: "You still haven't explained the ineffable aliveness of these organisms!"

There are obviously large families of differences and similarities in experience that are best ignored at this stage of inquiry—no one can get a good scientific handle on them yet. One can admit that there is a lot more to color experience, or any other domain of subjectivity, than we have yet accounted for without thereby endorsing the dubious doctrine that qualia are properties that elude objective science forever. But that doctrine is the standard destination of all the qualia arguments among philosophers.

MG: So what is the task of the future students of the problem of consciousness? What should be the content of their research? Is it to solve the brain mechanisms enabling, say problem solving, and along with that will come some deeper understanding of the old ineffable qualia?

DD: That's roughly right, in my opinion. Here is one place—not the only one, of course—where cognitive neuro-

scientists could take a hint from AI. The people in AI have almost never worried about consciousness as such, since it seemed obvious to them that if and when you ever got a system—an embodied robot, in the triumphal case—that actually could do all the things a person can do (it can reflect on its reflections about its recollections of its anticipations of its decisions, and so forth), the residual questions about consciousness would have fairly obvious answers. I have always thought they were right.

MG: The quip often heard about your book, *Consciousness Explained*, is that you explained it away. So, let me come at the problem from another angle. There can be little doubt most of our brain activity that enables us to do anything goes on outside the realm of our conscious experience. We hardly have access to the processes that allow us to be motoric, to create, to recall, and so on. We seem to know only about the products of these activities. What is that? What is it that is looking at all of these products?

DD: “We”? Who or what is this “we” you speak of who has or lacks access to various processes? A self is not a separate thing in the brain, with its own agenda and powers, which is made privy to some brain processes and not others. There is nothing that is, as you say, “looking at” all these products, though I agree that it is very hard to keep this strange fact in place as one thinks about what’s going on. The various *effects* of conscious access (or lack thereof) have to be shown to be the natural and indeed constitutive outcome of the activities and processes themselves, traced out through all their interactions. A sure sign of residual Cartesianism in any model is when it describes processes leading up eventually to some central transduction or threshold-crossing (or phase lock or induced synchrony), which is then declared, for rea-

sons good or bad, to ensure consciousness for the product in question. At any such moment we must go on and ask the embarrassing question: “And Then What Happens?” That is, what account does the model give of what is thereby enabled by this putative onset of “access”? Most models give no account at all. The task of the cognitive neuroscientist, however, is not just to explain how one’s favorite phenomena get all the way up to consciousness; to complete the task one has to explain what happens all the way *through* consciousness to eventual behavior (and behavioral dispositions, of course). Only then will we be able to see why and how the theory is a theory of consciousness at all.

The quip that my book ought to be titled *Consciousness Explained Away* is telling. Different readers no doubt have different grounds for saying it, but in any event it would perfectly express the attitude of one who had missed the whole point of the book—rather like somebody who might quip that Darwin’s theory of evolution by natural selection didn’t so much explain the design in the biosphere as explain it away. My theory of consciousness certainly doesn’t explain everything about consciousness that needs explaining, but at least it has the right overall shape: it undertakes to show how each feature that people have taken consciousness to exhibit is either the effect of some mechanism or mechanisms the operation of which can be understood without any tincture of consciousness, or else is the figment of an inflated or otherwise mistaken claim. I don’t see how any other *sort* of theory of consciousness could presume to have explained it. Has liquidity been explained away by the physicists because, in their final account, they don’t attribute liquidity to anything at the atomic level? The physicists have left out the wetness, and I’ve left out the qualia. On purpose.

MG: But in the case of physicists explaining away wetness, they can reconstruct every aspect of actual wetness from their molecular theory. They can show how surface tension necessarily creates drops, how the rolling and tumbling among molecules of a liquid state allow it to pour and assume the shape of a container, and so on. But in the case of consciousness, can your theory actually show mechanically why my pain “hurts” me (as opposed to merely changing my goals and behavior) and that apples actually “look red” to me (as opposed to merely contrasting with leaves and reminding me of firetrucks)?

DD: You are certainly right to stress that the effects still in need of explanation are many, but there is a fatal—and common—mistake to avoid here: arriving at the “conclusion” that after “all” the effects of this sort are explained, there will be some inexplicable residue. How do some people reach this imagined conclusion? By imagining themselves to engage in a process of something like subtraction: “Here am I, looking at the apple, and reflecting on how wonderfully red it appears. Now I subtract my reflections, my dispositions, my changes in mood, my memories, my . . . and I ask: ‘what’s left?’ and I ‘see’ that there is still something left over: the very intrinsic redness of it all!” That is not an argument; you couldn’t prove anything with such an exercise of the imagination, if only because there’s really no way you can prevent the very items you take yourself to have subtracted away from somehow returning surreptitiously to fuel your sense that something is still there.

Compare it to the naive but strangely compelling attitude some people have toward dollars, encapsulated in the American tourist’s query: “What does it cost in *real* money?” Such a person finds it easy to believe that marks and francs and pounds and yen have value only in virtue of their exchange

rate with dollars, but they persist in thinking that dollars are different; dollars have *real* value, *intrinsic* value! These people find it very hard to believe that there isn’t “something left over” when they’ve subtracted all the merely dispositional features of dollars—their instrumental value in exchange for goods, services, and other currencies. They are wrong, of course. I am claiming that the hardcore qualophiles are making the same sort of mistake.

MG: So this brings us to your own strategy of discovering new insights in the stuff of conscious experience. Are you not trying to build a cognitive/conscious agent at MIT? Tell us about that project and, in particular, speak to the point of Searle and others that building agency out of anything save biological material is a doomed enterprise.

DD: Cog, undoubtedly the most ambitious, most humanoid robot yet attempted, is being designed and built at the AI Lab at MIT, by a team of graduate students under the direction of Rodney Brooks and Lynn Andrea Stein. I am playing an advisory role on the team, and, in the process, learning all my heart desires about the immense technical difficulties of building actual robots.

Cog is to have an extended “infancy,” not growing in size, but developing many of the competences that human infants develop, from thousands of hours of embodied “experience” in the real world. Cog is adult-size, with a movable torso, head, and arms, but lacking legs. Cog is bolted at the “hips” to a fixed pedestal, which solves the problem of providing huge amounts of electrical power and multifarious connections to Cog’s massively parallel brain, which is telephone-booth-sized, without a cumbersome trailing umbilical cable. Cog’s fingers, hands, and arms have approximately the same amount of “give” as their human counterparts, and Cog’s

eyes saccade at near-human rates (3, not 4 or 5, saccades a second, with comparable speed of saccading and dwell-time). Cog's eyes are composed of two tiny TV cameras, a high-resolution foveal camera mounted on top of a wide-angle parafoveal camera. Among the features of human vision that have to be modeled in Cog are the problems of integrating the VOR, head and skeletal motion in addition to eye movement, vergence control, motion detection, "pop-out" for various importance features, face-recognition, . . . the list keeps growing, of course. Achieving human-level hand-eye coordination is a central goal, but before that can be addressed, we have to ensure that Cog won't poke its eyes out with inadvertent motions of its arms! So a pain system, and innately "hard-wired" (actually software-controlled, of course) avoidance of such mischief is a high priority.

It is still too early to say just how far, and how fast, the Cog project will go, but at least the problems being addressed are real problems of real cognitive science, shockingly oversimplified from some perspectives—from the standard perspectives of functional neuroanatomy, for instance—but still orders of magnitude more realistic than other modeling efforts in AI. The Cog project is controversial among people working in AI, and some outspoken critics think it will come to much less than the fascinated public (and science journalists) expect, so much less that it is an unwise undertaking at this time. I disagree, but of course I am biased. For me, it is like being given Aladdin's lamp: with any luck, I will soon know whether some of my favorite inchoate ideas can be turned into working models, a task that is way beyond my own technical competence, but well within the range of the brilliant young people on this team.

One of my advisory roles is directing members of the team to crucial ideas, phenomena, and problems, from other areas

of cognitive science that they have not yet encountered on their own. They are primarily engineering students, but quick studies with voracious curiosity, undaunted by any technicalities. I mention this in particular, because any cognitive neuroscientists who have a burning conviction that Cog will never work without X (where X is something they know all about) are invited to try to convince the Cog team (through me) that they are right. In other words, short, argument-packed letters that begin, "If I were designing Cog's vision system [motor-control system, audition, memory, pain system], I'd make sure that it exploited . . ." will be carefully read. We don't think we already know all the answers about how to do it.

One thing we're sure about, though, is that John Searle's idea that what you call "biological material" is a necessity for agency (or consciousness) is a nonstarter. Oh, it might turn out, for largely boring reasons, that electric motors are such poor substitutes for muscles (made of organic polymers, artificial or natural), that any truly effective humanoid robot must have organic muscles. And I suppose it might turn out for similarly boring reasons that silicon chips, no matter how massively parallel, simply cannot do all the transformations (= computations) that the organic materials in our nervous system do, but if this turns out to be so, it would not be any confirmation of Searle's vision, since he explicitly detaches the "causal powers of the brain" that he is interested in from all such issues of real-time control. He concedes (perhaps unwisely) that a silicon brain could control a humanoid body exactly as well and as fast as an organic brain. If that is so, Cog can get by just fine with silicon chips, which is what we are gambling on.

MG: But even if qualia or subjective experience can be explained right out of science, aren't they ineliminable from

the very way we think about ourselves and each other, and especially from ethical thinking? The whole argument about animal rights has to do with whether the fish actually feels pain when it bites the hook, or just flops around reflexively. If I cut the cord to Cog, then I'd be guilty of vandalism if it didn't have any conscious experiences, but I'd be guilty of murder if it did. So it seems like the sense of consciousness you want to explain away really does make a difference!

DD: I agree that it is ethical considerations that make the question of pain, and hence consciousness, so important, and this is exactly why it is not just wrong but deeply immoral to mislocate the issue in doctrines that are systematically unconfirmable and undisconfirmable. If the question of whether the fish feels pain is declared to be unknowable in the limit of scientific inquiry then how on earth could the injunction not to cause unnecessary pain be so important? What is important can be observed, shared, noticed—if not yet, then by an extension of investigations we already know how to conduct. I think the idea that pain is, as it were, a morally important but nevertheless unmeasurable “quantity” is a pernicious oversimplification (as I argued in the section called “Minding and Mattering” at the end of *Consciousness Explained*). In the case of Cog, I agree entirely that the time may well come when our moral duties to Cog (and not merely to Cog's owners) become a very serious consideration, for exactly the same reasons they are a consideration for any experimenters working with animals (including human beings). There has already been considerable discussion about this among members of the Cog team and interested on-lookers. And let me end on a reassuring note: the errors will almost certainly be on the side of oversolicitousness. People—even the sophisticated technocrats who make robots—are amazingly easily moved to sympathy, empathy,

concern. A little “eye” contact is overwhelmingly moving. If Cog “works” at all, you can rest assured that Cog will have plenty of ardent guardians, eager to weigh Cog's own interests and needs in any decision making.

MG: So your position is there is really no conceivable argument against a functionalist view, given our knowledge and beliefs about the explanatory power of modern science?

DD: Oh, I'm sure we can *conceive* of arguments against functionalism; it's just that I haven't encountered any good ones yet. But who knows what argument will come along tomorrow? I certainly don't want to encourage neuroscientists to turn a deaf ear to philosophical arguments—open-minded skepticism seems to me to be the appropriate attitude.

MG: Well, laboratory scientists are always fascinated with philosophy and philosophers. One thing that always comes across is how trained and expert philosophers are in the art of argument and in the distinctions they insist on making. At the same time, sometimes it is felt philosophers and in particular the modern philosophers of mind stake out positions and then consider new data from their personal perspective, not with the aim of validating or invalidating their view but with seeing how to keep their view intact given the data. Now this is not an impudent charge. It is a reflection of the fact that since we are a light year or two away from truly understanding how the brain does its business, this is the only practical way to survive. Or would you reject this interpretation of current behavior?

DD: I see it a little differently. Scientists just as often as philosophers defend their positions until the last dog is hanged, and so they should. You don't abandon a promising theory in the face of a single unforeseen counter-instance if you can

think of a way to refine or adjust your theory. Human nature being what it is, however, we are often tempted to preface such a regrouping with “What I meant all along was . . .” instead of, “What I should have said was . . .” But philosophers are actually in a slightly different position from other theorists. We philosophers have a delicate balancing act to perform: as would-be analyzers of concepts, among the truths we strive to uncover are conceptual truths, and these shouldn’t be any more vulnerable to straightforward empirical disconfirmation (or confirmation) than their more obviously a priori brethren, mathematical truths. So it is entirely appropriate that we try to construct theories that leave most of the empirical options wide open—it is not our job to fill in all those details. So any time anything we say appears to be flatly at odds with some empirical discovery, something has to give. Most often, the right thing to do is to re-express the philosophical point in a way that shows that it was not foreclosing on the discovery after all. And almost as often, the nonphilosophical critics actually *have* misinterpreted the philosopher’s position, so a certain amount of “you’ve misunderstood me” is perfectly legitimate! Suppose a bridge collapses, and we confront the geometer who advised us on its construction: “We thought you said triangles were rigid figures!” we complain. “And so they are,” he replies, undaunted by the pile of twisted steel members. “These are *former* triangles.” What else should the geometer say—that triangles are *unusually* rigid, or that they are rigid unless undue strain is put on them? Those aren’t truths of geometry. Notice that truths of geometry do explain why bridges made of triangles are sturdier than bridges without them—and these explanations embody testable empirical predictions. The geometer isn’t copping out, and philosophers *need not* be copping out when they point to an escape hatch in their definitions.

MG: Finally, then, help us distinguish the major views of current philosophers of mind. You, as the supreme functionalist, hold that an artifact that was complex enough could have all the properties of consciousness. The Searle school would reject this and maintain that there is something special about neural tissue that makes it a necessary substrate or source of consciousness. And finally, the Churchland school maintains that in order to explain how brain processes are conscious processes, you have to descend to principles at the molecular level. Is that roughly right? Are there other contenders we should know about? Where does that leave the cognitive neuroscientists? Are we waiting for you to add to the debate or are you waiting for us?

DD: It’s interesting to see just how the philosophical disputes appear to you—and no doubt to your colleagues. Let me suggest a few revisions. In fact, I see myself in agreement with the Churchlands about everything except minor details, mainly of emphasis and method. Unlike them, I am simply agnostic about how deep into the particular details of neuroanatomy or neurochemistry we will have to go to get models that work—that can have the input-output functions required of minds. Even if we do have to go to the molecular level, I’ll still consider functionalism unscathed; it will just have turned out that there are many *less* ways to skin the cat than I had supposed! I recently conjectured in a playful spirit that it *might* even turn out that just as some of the microscopic endoparasites in our gut play a well-nigh ineliminable role in our digestion, so other macromolecular parasites in our nervous systems might be required for cognition! Unlikely, surely, but as a worst-case scenario, it shows that functionalism is not *committed* to any particular “high level” of modeling. If Penrose and Hameroff are right—and

I'll eat my hat if they are—functionalism will have to descend to the quantum level to find its proper footing. It turns out that you can make quite serviceable artificial hearts without copying organic hearts at even the level of gross anatomy: artificial brains will no doubt have to be a lot more like organic brains to do their stuff, but how much is still an open question.

The Churchlands think they know what the right level for modeling minds in brains is. They might be right, but I'll reserve judgment. Given their views, they have expected more radical conceptual revisions to arise from neuroscience—overthrowing or “eliminating” the categories of folk psychology—while I have stressed that folk psychology (such everyday categories as belief, expectation, intention, dreaming, pain) are so powerfully predictive and useful that they are here to stay. So I have sought a more indirect accommodation of these categories within neuroscience. In the end it is not so much a factual or even theoretical disagreement between us as a tactical one, parallel to the simpler question of whether physicists should say that they have an explanation of “centrifugal force” or an explanation of why there really isn't any such force at all. (I think my disagreement with Pat Churchland over “filling in” is largely due to her misunderstanding my position, but that has been partly my fault—one of those cases where I have in fact prefaced my rejoinder with, “What I should have said was . . .”) The main point of theoretical agreement between us is that what happens in the brain does *not* map neatly onto everyday notions of the mind (for instance, there is no Cartesian Theater in the brain, but it sure seems as if there is!), so materialism is a harder, more radical doctrine than some have thought.

Searle is not even in the same discussion. He claims that organic brains are required to “produce” consciousness—at

one point he actually said brains “secrete” consciousness, as if it were some sort of magical goo—but since this is for him just an article of faith with no details, no models, no explanatory power, no predictions, it is hard to know what to say in response. Given the peculiar way he divorces his favored “causal powers” of brains from their control powers—the powers that permit them to accomplish discrimination and perception, underlie memory, guide behavior—his doctrine is conveniently untestable, now and forever. He and his followers do not shrink from this implication—they embrace it! To me, this is an unvarnished *reductio ad absurdum* of his position, and I marvel that anybody takes it seriously. Some people just love an insolvable mystery, I guess.

MG: Thank you.