

## What Consciousness Is Not

*Raymond Tallis*

With the publication of his 1996 book *The Conscious Mind: In Search of a Fundamental Theory*, David Chalmers established himself as one of the most assiduous, honest, imaginative, and talented thinkers working in the vast and overpopulated field of the philosophy of mind. In that tome, Chalmers did not avoid the abstruse and the technical where they were unavoidable, and only intermittently lost touch with the mysteries that strike us all when we think about consciousness. And for the most part, despite the difficulties, he also managed to explain his inquiries with admirable clarity; in this respect, he came across like the philosopher John Searle, only less combative, less sure of himself, and less liable to brush aside or overlook the true problems of consciousness. (Searle, incidentally, launched a savage attack on Chalmers's book.) If Chalmers's scrupulousness and attention to contrary views made his arguments long—sometimes wearily long—this was an indirect tribute to his seriousness of purpose.

The opposite of a sophist, Chalmers, then a professor at UC Santa Cruz and now a professor at Australian National University and New York University, seemed in the book that launched him to prominence like someone who really wanted to advance our understanding, and his own, rather than simply to win adherents to a position. Nor was he a mere pedant. He proved more willing than many others in the field to experiment with views that were—and indeed still very much are—radically at odds with philosophical orthodoxies. Those views made Chalmers a target for the mockery of that orthodoxy—hard-line materialists such as Searle and Daniel Dennett, who, notwithstanding the heavy trade of insults that has passed between those two bruisers over the years, are both dedicated to the fundamental notion that the mental and the neural are one, and that the conscious mind is a solely biochemical phenomenon.

Chalmers's latest book, *The Character of Consciousness* (2010), at least begins well, and there is much in it to recommend. The early chapters are a significant contribution to clarifying and deepening the fundamental questions of the philosophy of mind:

---

*Raymond Tallis, a New Atlantis contributing editor, is emeritus professor of geriatric medicine at the University of Manchester, United Kingdom, and the author, most recently, of Aping Mankind: Neuromania, Darwinitis and the Misrepresentation of Humanity (Acumen, 2011).*

---

What is consciousness? How can it be explained? Can there be a science of consciousness? What is the neural basis of consciousness? What is the place of consciousness in nature? Is consciousness physical or nonphysical? How do we know about consciousness? How do we think about consciousness? What are the contents of consciousness? How does consciousness relate to the external world? What is the unity of consciousness?

The multiplicity of these questions is to be entirely expected, given that consciousness is, as Chalmers puts it,

an extraordinary and multifaceted phenomenon whose character can be approached from many different directions. It has a phenomenological and a neurobiological character. It has a metaphysical and an epistemological character. It has a perceptual and a cognitive character. It has a unified and a differentiated character.

And that's just for starters. The mystery of consciousness is a network of mysteries, touching on the mystery of ourselves, the mystery of the intrinsic nature (if any) of the non-conscious world, and the mystery of our *knowledge* of ourselves, the natural world, and the human world atop it. If there is such a thing as a First Philosophy, the philosophy of the conscious mind is it. It is the ground in which every other branch of philosophy takes root.

Considering the profound importance of these questions, Chalmers's latest book, *The Character of Consciousness*, ultimately turns out to be a disappointing sequel, especially given his track record of taking on the conventional wisdom that the answers to these questions are likely to defy. But it is worth considering this book at some length; for given David Chalmers's distinctive sobriety and thoughtfulness among a field of philosophers committed to reducing its chosen subject nearly out of existence, it is striking how much his work still falls prey to the same fundamental errors. The book will thus serve as an instructive case study not only in how befuddling are questions about the mind, but in how stuck is the philosophical rudder of the prominent thinkers who study it, and how adrift they have floated.

*The Character of Consciousness*

By David J. Chalmers

Oxford ~ 2010 ~ 624 pp.

\$29.95 (paper)

*The Conscious Mind:*

*In Search of a Fundamental Theory*

By David J. Chalmers

Oxford ~ 1997 ~ 432 pp.

\$24.99 (paper)



---

FALL 2011 ~ 67

## The ‘Easy’ and the ‘Hard’ Problems of Consciousness

*“What is the answer?” she asked, and when no answer came she laughed and said: “Then, what is the question?”*

—Reputedly the last words of Gertrude Stein

It was as the clarifier of questions that David Chalmers made his initial reputation. It was he who first proposed the now-standard distinction between the “easy” and the “hard” problems of consciousness. The easy problems are “those that seem directly susceptible to the standard methods of cognitive science, whereby a phenomenon is explained in terms of computational or neural mechanisms,” while the paradigmatic hard problem is “the problem of *experience*.” An organism possesses the trait of experience when we can say that it is *like* something to *be* that organism, as Thomas Nagel put it. This is true most notably and most elaborately of organisms like you and me; it is also probably true of most vertebrate animals, but probably not true of any plants.

Philosophers of mind call specific kinds of experience *qualia*. These are the most basic elements that underpin the hard problem of consciousness, and include, for example, the feeling of warmth, the taste of wine, the way the color red looks, the sensation of pain, and the sound of birdsong. It is these that fill conscious life. Moreover, in the case of human beings, an experience is always had *by* someone (and *only* by that someone): it belongs to first-person being, and so lies beyond the reach of third-person (or no-person) objective science. Anyone can witness and so equally partake in the fact that some physical event occurred; but an experience occurs only for some individual person. (Consider the difference between watching a man stub his toe on a table and experiencing the resultant pain.)

The distinction between hard and easy problems is useful—even profound. But it has been deeply problematic since Chalmers first proposed it, for the place where he draws the boundary between the two yields too much to those who believe neural and computational science can fully explain consciousness. He is too ready to consider parts of the mind that do not entirely belong to the physical world as being no more than physical way stations in the causal chain between sensory inputs and behavioral outputs.

Chalmers’s list of “easy” problems includes: our abilities to describe our mental states, to focus our attention, and deliberately to control our behavior; how cognitive systems acquire and integrate information; and the difference between wakefulness and sleep. This should not be allowed

to pass on the nod. One would have thought that information, attention, deliberation about controlling behavior, and wakefulness are things about which we can answer the question, “what is it like?”—so, for that matter, is dream-filled sleep, and so are mental states of just about any kind. Indeed, if these mental features did not *feel like* anything, they would not be what they are supposed to be (and a difficult, indeed paradoxical, set of questions would still remain about why they at least *seem* to feel like something). Chalmers’s failure to see this fatally damages the inquiry that follows.

Even the first and seemingly least contentious item in his list—in his words, “the ability to discriminate, categorize, and react to environmental stimuli”—should put us on red alert. Though the notions of “stimuli,” “environment,” and “categories” can in some sense be used in our descriptions of the behavior of unconscious organisms such as bacteria, these terms are applied by extrapolating from organisms with experiences, and so ultimately owe their distinctive meaning to organisms that are conscious in the “hard” sense. By saying that so much of consciousness is amenable to a solely neural-computational explanation, Chalmers gives too much ground to those philosophers who believe there is no fundamental difference between a conscious organism and a mechanism, because they regard the conscious mind as no more than a machine for linking environmental inputs to behavioral outputs in the most effective way. This concession makes Chalmers’s fundamental position that “consciousness is not physical” more vulnerable than it need be. And it opens him to the accusation of inconsistency from committed physicalists, for whom experiences are illusions—merely ontological spooks left over from “folk psychology.”

It is easy to see what unites the elements of the mind corralled together in the pen of “easy problems”: they all have direct behavioral correlates. They are usually associated with events in the physical world, observable by anyone, and so at least some part of them can be described as purely physical. And this, for some, means that these “easy” problems can be entirely reduced and translated into their behavioral correlates. The behavior associated with, say, waving your hand could be simulated by a zombie—a hypothetical being, often discussed by philosophers of mind, that *acts* like a conscious person but is not conscious. Because the actual *experience* of waving is in theory not required for you to wave your hand, or (more precisely) for waving to occur, it must be regarded as a sort of accidental add-on: even if the experience exists, it does not actually *cause* the behavior of moving your hand, but is rather a sort of bystander to the event.

---

Of course, even from the evolutionary perspective adopted by most of those who reduce the mind to a mere machine for transforming sensory inputs to behavioral outputs, there are near-fatal difficulties posed by the idea that experience is, as philosophers of mind call it, “epiphenomenal”: how could a trait that is incapable of affecting an organism’s behavior, and so its ability to survive and reproduce, be effective at propagating itself through evolution? (Why would evolution continue the existence of pain if your body could simply withdraw your hand automatically without it?) Moreover, even if pain *were* epiphenomenal, the experience would nonetheless still exist, and would need to be accounted for and explained. The supposedly “easy” problems, no matter how one comes at them, still have irreducible, “hard” elements. This fundamental error plagues the ideas of thinkers who are committed to a reductive account of the mind, and Chalmers’s work, despite his reputation for pushing back against that trend, falls prey to the same mistake.

### **Consciousness and its Place in Nature**

From the outset, there has been another strange consequence of Chalmers’s extension of the “easy” problems of consciousness to encompass activity that would seem to be inseparable from experience: the idea that consciousness extends in some sense to *the entire universe*, or that all matter is conscious. Known as *panpsychism*, it was this idea that earned Chalmers the scorn of John Searle. We can see what this means, and why Chalmers adheres to it, by examining his use of the word “information.”

One might be surprised to learn that anything involving “information” could literally be located outside of experiential consciousness. (The use of the word to refer to the contents of computers, books, and so forth is a matter of proxy: these objects store the *products* of a conscious person or persons, and only become “information” when they are translated back into the mind of another conscious person or persons. But as long as it is being stored in or shuffled between them, the content of computers remains “information” only in this indirect, honorary, or metaphorical sense.) Far from being eccentric, however, the view that stuff outside of consciousness is “information” is entirely orthodox. The use of the terms “information” and “information processing” to refer to events taking place in the brain is the mainstay of much cognitive science; the soundness of this idea is the central assumption of the Computational Theory of Mind that dominated cognitive science for nearly half a century. “Information” is used not in the sense in which you and I use it, to mean something that

---

is conveyed to me by someone or something, and of which I am conscious; rather, it is merely about the relationship between inputs and outputs in any system. Once this is accepted, we soon reach the position that Chalmers entertained in *The Conscious Mind*:

wherever there is a causal interaction, there is information, and wherever there is information, there is experience. One can find information states in a rock—when it expands and contracts, for example—or even in the different states of an electron. So...there will be experience associated with a rock or an electron.... It may be better to say that a rock *contains* systems that are conscious.

In light of this redefining of information, it should come as no surprise that in his new book, Chalmers argues that “we have good reason to suppose that consciousness has a fundamental place in nature” and that “consciousness and physical reality are deeply intertwined.”

Chalmers arrives at this position in part by a process of elimination of other ways of seeing the place of mind in nature. He first points to the problems with various kinds of reductive materialism, which holds that the physical realm and the phenomenal realm (that is, the realm of sensations, perceptions, and appearances) are actually the same things. (Some kinds of reductive materialism at least argue that there is a difference in how we *know* about the physical and the phenomenal, although one variety believes that this difference can be eliminated, so that we could view our own experiences as purely biochemical events.) Chalmers’s commitment to acknowledging the hard problem of “what-it-is-like” experience means he cannot hold to any form of reductive materialism. He then turns to dualism. One kind of dualism holds that physical and phenomenal events are distinct kinds of things, and that each affects the other: physical events can cause phenomenal events, and vice versa. This position, known as “interactionism,” is roughly the same as Cartesian dualism and inherits all its well-known problems. Then there is a kind of dualism which also holds that physical and phenomenal events are distinct kinds of things but the causation goes only one way, so that phenomenal events are *caused by* physical events but cannot in turn cause or affect physical events. As discussed earlier, this is the idea of “epiphenomenalism,” and it, along with interactionism, is beset by insoluble problems. According to Chalmers, this leaves only the idea that the intrinsic properties of the physical world are all also inherently phenomenal properties.

Surprisingly, Chalmers is not alone in giving serious consideration to this last position, which he calls “type-F monism”—monism being the

---

idea that there is only one kind of thing, as opposed to dualism, which asserts two. Indeed, the idea is gaining adherents: leading philosophers of mind, such as Michael Lockwood, Donald Griffin, and Galen Strawson, have advanced similar views. Type-F monism (so called because Chalmers dismisses approaches he labels A through E) has its immediate ancestry in Bertrand Russell's monism in *The Analysis of Matter* (1927), though it was first expounded by Spinoza in the seventeenth century. For Russell, the idea came from his observation that physical theories characterize physical things only in terms of how they relate to other physical things, but do not tell us anything about the intrinsic properties of those things—which leaves open the possibility that some of the properties of physical things might be phenomenal.

There are, however, many problems with type-F monism, which is a form of panpsychism. Most concerningly, it would seem to spread consciousness, or something like it, too far: it would *feel like* something to be an electron, as Chalmers hinted in his earlier book. This was the view mocked by John Searle; it reminds also of the Victorian novelist Samuel Butler's tongue-in-cheek line that "Even a potato in a dark cellar has a certain low cunning about him which serves him in excellent stead"—but the earnest claim about an experiencing electron seems even more vulnerable than the teasing one about the potato.

Chalmers tries to get around this by suggesting that there may be something more fundamental than consciousness—some "protophenomenal" property, as he calls it, which is necessary for phenomenal experiences, but not sufficient. It is *this* property that all matter possesses, and that is universally present in the physical world. Yet this seems a not especially helpful theory: matter itself already qualifies as the mysterious thing that is necessary but not sufficient for consciousness—but this insight is the beginning and not the end of the problem. Since we already well know that matter has the latent potential to become conscious, how does the notion of a protophenomenal property make the situation any easier to grasp than trying to understand consciousness in terms of what we already know about matter itself? Unfortunately, Chalmers simply moves the problem on: What is the difference between the protophenomenal properties of a non-sentient pebble and the phenomenal experiences of a frog, or the experiences of a fully sentient and thought-filled human being? What is it that enables the merely *proto*-phenomenal properties that supposedly pervade the air around us to become *fully* phenomenal properties when we breathe that air in and it becomes part of our brains?

---

With characteristic honesty, Chalmers admits the difficulty of answering questions like these:

Our phenomenology has a rich and specific structure. It is unified, bounded, and differentiated into many different aspects but with an underlying homogeneity to many of the aspects, and it appears to have a single subject of experience. It is not easy to see how a distribution of a large number of individual microphysical systems, each with its own protophenomenal properties, could somehow add up to this subject of experience. . . . Should one not expect something more like a disunified, jagged collection of phenomenal spikes?

The answer to this question is surely yes. Indeed, it is just about impossible to see how a distinct self such as “David Chalmers” could have been constructed out of the mere protophenomenal twinkling of the material world. Even less can one understand how that twinkling could in Chalmers distinguish between *itself*—a set of neural discharges, somehow aware of itself—and the rest of the world. His conclusion therefore seems to be more than a little optimistic: “Overall, type-F monism promises a deeply integrated and elegant view of nature. . . . [It] is likely to provide fertile grounds for further investigation, and it may ultimately provide the best integration of the physical and phenomenal within the natural world.” Fascinating and remarkable as is Chalmers’s long chapter on the metaphysics and ontology of consciousness, one feels compelled to draw the opposite conclusion: I wouldn’t start from here.

### **Neural Correlates of Consciousness**

One of the most impressive parts of *The Character of Consciousness* is Chalmers’s investigation of the relationship between neuroscience and the philosophy of mind. In recent decades, neuroscience has increasingly focused on identifying so-called “neural correlates of consciousness” and describing their characteristics. Hence the rash in the media of images from brain-scanning devices, that show parts of the brain “lighting up” as a purported explanation for various phenomena of the mind.

There have been various theories proposed as to how neural activity might become consciousness. One key fact, which is embarrassing for mind-brain identity theorists, is that the overwhelming majority of neural activity—in both the brain and the spinal cord—is not correlated with awareness of any sort. For some, this means that the answer is location: neural activity in the cerebral cortex is consciousness and neural activity in

---



the spinal cord or cerebellum is not. Yet location hardly seems an adequate explanation of how some nerve impulses get upgraded from mere biophysical events to bits of awareness. *What* is it about the cerebral cortex that is special? What qualities does it, or whatever part of the brain underlies awareness, possess that the others do not? What *physical* qualities could account for the supposed difference between awareness-causing neural activity in the cerebral cortex and neural activities elsewhere in the nervous system, which even hard-line reductionists agree are not conscious?

Various alternative theories appeal to neurobiological properties that are less anatomically localized. These include “systems,” such as the one emphasized most recently by Gerald Edelman, in which consciousness arises from “loops” of activity between the thalamus and the cortex. Similarly, Francis Crick and Christian Koch speculated that consciousness might involve a particular sort of cell throughout the cerebral cortex, which has “a unique combination of molecular, biophysical, pharmacological and anatomical properties.” Other approaches focus more on what the neurons are up to than where they are—their patterns, their intensity, their frequency, the extent to which they are synchronous, and so on.

But none of these characteristics seems likely to deliver the difference between neural activity that is and is not associated with consciousness, not the least because they all aim to narrow down a phenomenon that is inherently multifaceted. And the approach faces other inherent limitations. For a start, as Chalmers points out, correlation is not causation: even if one identifies some neural feature correlated with consciousness (say, by stimulating a part of the brain and having the subject report being aware of some mental state), it does not follow that this neural feature is solely or mainly dedicated to consciousness. More to the point, even if some of these phenomena do turn out to be truly and uniquely causative of consciousness, none of them would enable us to get a handle on the “hard” questions. As Chalmers candidly points out, “why should [some particular neural feature] give rise to conscious experience? As always, this bridging question is unanswered.”

### **The Unity of Consciousness**

Any theory of mind will have to address the particularly intractable feature of consciousness that Chalmers encountered in advancing his theory of type-F monism: its unity. The many experiences that we have at any one time—sights, sounds, smells, and other sensations, as well as memories, thoughts, and emotions—are in some sense unified. They belong to a sense

---

of *me-here-now*, to what we might call *co-consciousness*—consciousness of many things all at once. This poses a radical challenge to any putative neuroscience of consciousness. The different contents of consciousness are supposed to be kept apart by existing in different parts of the brain; yet they are also required to come together somehow. And however this convergence is accomplished—say, by merging the pathways between the different parts of the brain—those contents would always seem in the act of becoming unified to lose their distinction, becoming boiled down to some unholy soup of undifferentiated awareness. The mystery is that the field of consciousness *is* unified while still retaining the distinctiveness of its contents.

This is the so-called “binding problem,” and there have been many attempts to find a solution to it. These attempts mostly depend on the idea that certain physical properties *common* to large swathes of the brain can *bring together* activity scattered across different regions: all of the different regions will be activated at once in the moment of consciousness, but without losing their spatial separateness. Proposed candidates for the special properties have included electromagnetic fields, quantum coherence, and synchronous electric oscillations in large sections of the cerebral cortex. But all of these candidates fail for the simple reason that they rely on *objective*, or externally observed, unity being translated into *subjective*, or immediately experienced, unity, with no reason offered for why this should be enough. If we accept that physical unity creates experiential unity, then the anatomical unity of the brain or, indeed, the body, should be considered of equal power to explain the unity of consciousness. But this suggestion makes clear that the theory would deliver too much, conflicting unification.

Chalmers devotes an entire chapter (in which he is joined by coauthor and Oxford professor Tim Bayne) to this mystery. He begins, usefully, by teasing out the different aspects of conscious unity. First, there is the unity of *elements* of consciousness that are all focused on the same object, as when I look at an item and see that it is red, it is a book, it is oblong, and so forth. Then there is the unity of two items in my visual field that are seen to belong to the same space, as in a car that is near to a tree. There is also “subject unity,” where two elements of consciousness are experienced as both belonging to the self at the same time. These kinds of unity can be further distinguished, between phenomenal unity, where two conscious states are experienced at once (so that it is *like* something to experience both at once), and “access unity,” where the two states can be accessed at once for verbal report, reasoning, and deliberate control of behavior. And at the most fundamental level, there is “subsumptive unity,” in which all

---

of the self's experiences are subsumed into a unified conscious field. This unified field is not simply the conjunction of all of the self's conscious states, but *is* a conscious state in and of itself.

The main aim of Chalmers in establishing the inescapable reality of conscious unity is that it is incompatible with some theories of consciousness that he wishes to discard. Among the problems besetting the search for neural correlates of consciousness is the unavailability—indeed, impossibility—of a “consciousness meter” which would allow us to directly detect and record consciousness. Instead, we can only infer the presence of consciousness through behavior, most obviously the behavior of someone verbally reporting that he is having a particular experience. This limitation has inspired one of the most ludicrous theories of consciousness: the Higher Order Thought (HOT) theory, advanced by City University of New York professor David Rosenthal. According to HOT theory, a mental state is conscious if and only if a subject is articulately aware of it. For example, I am phenomenally conscious of seeing red when I am having the thought that I am seeing it. So consciousness arises when thought lights up unconscious sensation.

HOT theory seems to turn everything upside down: surely thoughts *about* something would seem to depend upon *having* the conscious experience of that thing in the first place. Further, the having of thoughts requires at least some linguistic, or pre- or proto-linguistic, mode of communicating with oneself. HOT theory is much like the idea that in order to be conscious you have to be self-conscious, or conscious of being conscious—but it runs even more quickly into an infinite regress.

Moreover, Chalmers rightly argues that HOT theory provides no way to understand phenomenal unity. While two conscious mental states A and B could be unified by the subject having the thought “I am experiencing A and B,” this would be unusual. Just to become conscious of an object would require becoming conscious of all of its characteristics; this would be impossibly burdensome, as the thought corresponding to *what I am seeing* now would be endless. The prospect becomes even more mind-boggling when one attempts to unify the entirety of a conscious field—to have a *thought* about everything at once of which one is conscious (including the thoughts themselves).

Chalmers also uses the evident unity of consciousness to critique representationalism, the hugely popular theory that a mental state is conscious when it *represents* some object in the world, and when that representation plays a functional role in how the mind creates “outputs” in the form of behaviors. But if behavioral output is the direct product of

---

consciousness, then the unity of consciousness would require a unified behavioral output. This is clearly not true: our behavior is not unified in the way our consciousness is. I do not, for example, behave in response to an entire visual field, only to something in it judged to have salience.

It is difficult to understand why anyone should want to defend either of these theories in the first place. Representationalism seems as vulnerable as HOT theory. A *re*-presentation requires some primary *presentation*, and it is presentations that are the fundamental stuff of consciousness: things in the world, and in our own bodies, are *presented to* our consciousness, and are *present in* it; and they must be before they can be reflected upon or re-presented. This is true just as much as visual objects are the necessary precursors of *images* of objects, of reflections of what is there. Representationalism and HOT theory both place the derived contents of consciousness before the primary, the mediated before the immediate. This inversion is attractive because it is rooted in a longstanding metaphor that says the mind is a place in which the world is mirrored. The fact that one is conscious *of* objects out in the world seems to be explained by the idea that consciousness shares properties with those objects: the object is replicated in the mind in patterns of neural activity that are somehow isomorphic with the object. This notion traces back to Aristotle's claim that perception takes the form of an object without its content, rather as a mirror image extracts the appearance of the object while leaving the object itself intact.

In other sections of the book, Chalmers expresses some sympathy for representationalism, giving it more consideration than it deserves, and so we will return to it in due course. But it is easier to understand why he chooses to attack it, along with HOT theory. For these theories aim to elude the hard problem by moving phenomenal consciousness to the territory that Chalmers has conceded to the easy problems. HOT theory reduces consciousness to "access" consciousness—to that which I can describe and report. And representationalism moves phenomenal consciousness even further in that direction: the representation is *caused* by that which it represents, and it is to be characterized in terms of its functional role. In short, representationalism fits within the materialist notion that consciousness is fully wired into the causally closed material world, and the mind is simply a way station between sensory inputs and behavioral outputs.

### **Could Philosophers Be Zombies?**

Since University of Nottingham professor Robert Kirk first released them in 1974 to roam the dark alleys of philosophical letters, zombies—beings

---

that are physically identical to us in every way, and that behave exactly as we do, but lack consciousness—have been, so to speak, gnawing at the brains of philosophers. As we will see, the zombie thought experiment is less interesting in and of itself than for what it reveals about how prone philosophers of mind are to depart from the actual substantive questions into realms of technical minutiae and hypothetical fancy.

Contemplating these hypothetical beings raises some key questions for the philosophy of mind: Does the fact that you can (in principle) know everything about the physical processes of an organism without being certain that it is conscious mean that consciousness is something real and additional to physical processes? If knowing that a being is conscious means knowing more than the details of its physical composition, does that have metaphysical implications?

Ever so slightly closer to the ground, if zombies are possible, would we not have to conclude that consciousness plays no role in our behavior? On one view, this conclusion plays into the idea that, although our actions *seem* to be those of our free mind, they are really just dumb physical processes. This in turn suggests that consciousness is superfluous, if not outright illusory. On a less radical view, the possibility of zombies would mean that consciousness is not reducible to the physical, and so it is something over and above the physical.

It is the latter conclusion that Chalmers is particularly interested in. More specifically, he says, materialism is the idea that physical theory can account for everything that exists; but if zombies are possible, then something exists that physical theory does *not* account for, so materialism would be false. This may seem a small victory: even in this case, physical theory would still account for everything that *physically* exists, including all of our actions and every physical event that occurs—which would leave materialism relatively unscathed and consciousness again reduced to a powerless bystander—but this is nonetheless what Chalmers is after.

These conclusions would be relatively straightforward if zombies were possible. The key question then is whether zombies really *are* possible—and, perhaps surprisingly to non-philosophers, this is where things start to get confusing. For Chalmers wishes to argue that zombies are possible based on the fact that they are *conceivable*. This, one might imagine, is no small task. Is it valid to draw ontological conclusions from what we can think about?

“Conceivability” is a slippery notion, whose boundaries are determined both logically and psychologically. In order to set aside psychological limits to conceivability, Chalmers introduces the notion of “ideal

---

conceivability,” in which what is conceived of cannot be ruled out even on ideal rational reflection by a person without psychological limitation. If this gives the impression of circularity, nothing that follows dispels this impression. Nevertheless, his discussion of the relationship between logical, physical, and metaphysical possibility, conducted in a whispering gallery of objections and rejoinders from a large cast of thinkers, is worth following. For not only is this sort of discussion the thing that is crucially needed but typically missing from those thought experiments of which philosophers of mind are so fond, it also ends up becoming a cautionary tale for how prone those experiments are to spiraling away from the realities they are supposed to help account for.

Numerous arguments have been launched against the premise that conceivability implies possibility. Chalmers deals with fourteen of them, the most important of which derive from Princeton philosopher Saul Kripke’s argument that there are many true statements of identity that are both *necessarily* true and true as a matter of empirical discovery rather than logical necessity (that is, true *a posteriori*). For example, water and H<sub>2</sub>O are identical; this is a necessary truth of physics, but it was an empirical discovery—there was a point in time at which it was not known. So we do indeed have an example of something that is *conceivable* but is not possible—namely, it is conceivable that water and H<sub>2</sub>O are not identical, but it is not possible. Therefore we cannot agree that conceivability, or logical possibility, implies metaphysical possibility.

Chalmers deals with this by appealing to a response to Kripke made by M.I.T. philosopher Robert Stalnaker, who suggests that the statement in question actually expresses not one but *two* propositions. “Water” refers first of all to the *sensory* thing we experience; and this “watery stuff” could have turned out to have been made of H<sub>3</sub>O or XYZ rather than H<sub>2</sub>O. (To put it another way, we can imagine some possible world where what we experience as “water” turned out to have been made of something different.) In this sense, the proposition that water and H<sub>2</sub>O are identical is only contingently, not necessarily, true. But “water” also refers to whatever the word *actually* picks out in this world—our physical world—and *this* connection is necessary, since the physical stuff known as water is identical to the physical stuff known as H<sub>2</sub>O. Based on this distinction, Chalmers claims that it really is metaphysically possible that the sensory stuff we call “water” could be not identical with H<sub>2</sub>O. And so, it seems, the notion that conceivability entails possibility is safe, and the zombie thought experiment does after all show that consciousness is not reducible to the physical properties of an organism.

---

The purported water/H<sub>2</sub>O relationship has also been invoked as an analogy by others in explaining certain features of consciousness, such as the relationship between perception and neural activity: perception and neural activity, this explanation says, are identical, but they may not *seem* identical because they are different aspects of the same thing. An initial problem with this idea is that two different aspects of the same thing are *not* in themselves identical: the front of a house is different from the back of a house, even if they are both views on the same house. This is related to a point I have elaborated elsewhere (in these pages, “What Neuroscience Cannot Tell Us About Ourselves” [Fall 2010], and more recently in *Aping Mankind: Neuromania, Darwinitis and the Misrepresentation of Humanity* [Acumen, 2011])—which is that physical events do not *have* intrinsic appearance, much less something corresponding to the contrast between appearances and what they refer to. Yet this undermines the explanation that perception and neural activity are two different aspects of the same thing, for “aspects” are a *kind* of appearance. That is, aspects are *already* a content of consciousness, and so cannot be appealed to in order to *explain* consciousness. More specifically, it solves nothing to appeal to *different ways of perceiving* as a way of explaining perception.

Though Chalmers does not himself directly invoke the water/H<sub>2</sub>O relationship to explain consciousness, it should be clear at this point that he has really gone a long way around to wind up right back at the questions he started with. For as we have seen, the water/H<sub>2</sub>O explanation is *already* laden with appearances, yet it is just things like appearances that the example is invoked to explain. And not only is “water” an appearance, but so, too, is “H<sub>2</sub>O”—albeit a mediated appearance, removed through theory and measurement. Both are ways of understanding and looking at the same physical stuff, and so are different appearances of that same physical stuff. Of course, it is always conceivable that two appearances could refer to two different actual things; but in the case of “water” and “H<sub>2</sub>O,” it is not metaphysically possible, because they both refer to the same thing. So we do, after all, have an example of something that is conceivable but not metaphysically possible—and Chalmers’s case for the possibility of zombies remains unjustified, and cannot be used to argue that the mind is irreducible to the physical.

Yet the reason Chalmers’s logic fails has to do precisely with the nature of appearances, and with how they pervade our theories and attempts to understand. Specifically, appearances are distinct from the things they are appearances of and different appearances are distinct from each other. Both of these facts mean that an appearance is inherently distinct from

---

the neural firings it is correlated with: the two are not identical. And all of these facts are evidence enough that consciousness is not reducible to the physical, without having to expend dozens of pages attempting to demonstrate that zombies are possible.

The outcome of the conceivability-possibility discussion would thus seem to be finely balanced. The conceivability of zombies proves nothing (not least because a richer conception of consciousness, which acknowledges and accounts for its role in behavior, would not permit zombies to be ideally conceived of at all). Unless, of course, we notice the obvious: that a world in which we can intuitively understand the difference between zombies and ourselves is probably not ontologically unitary—not made of just one kind of thing. Our ability to conceive of the zombie thought experiment may not indicate that zombies are possible, but it does indicate that, whatever its role, consciousness is something over and above its involvement in our behavior, and there is a real distance between a conscious being and the material world it inhabits.

### **Concepts and Contents of Consciousness**

Perhaps not surprisingly, it is around about here that Chalmers seems to lose touch with the problems he set out so clearly in his earlier book and in the opening of the present one. The forty pages he (joined in this chapter by coauthor and fellow Australian National University philosopher Frank Jackson) devotes to the question of whether conceptual analysis is required to justify reducing phenomenal truths to physical truths seem to go round in circles. This is in part due to his continuing to run with the assumption that, even though it seems a matter of empirical fact whether or not consciousness is identical with neural activity, it seems also to be amenable to investigation by purely logical argument. He justifies the appeal to conceptual analysis on the grounds that it is supposed to be of particular relevance as to whether the epistemic gap (that is, the gap in how we get to know them) between the domain of physical processes and the domain of consciousness implies an ontological gap (that is, a gap in what they really are).

Chalmers's case is not persuasive, and the intellectual effort required to follow a thread that rarely rises above the technical is not rewarded by proportionate enlightenment. As we saw already when entangled in the technical weeds of the zombie thought experiment, the mystery of consciousness drains away into gray dunes of endless argument and counterargument, as we progress with all the celerity of a lichen growing over a rock

---



to the opaque conclusion that “if the phenomenal is reductively explainable in terms of the physical, then there is an *a priori* entailment from physical truths, indexical truths, and a that’s-all statement to phenomenal truths.”

It is therefore a relief (though only partial) when Chalmers turns to examine consciousness itself. He begins by making a sharp distinction between two sorts of mental states: experiences and beliefs. Experiences are “paradigmatically phenomenal” and beliefs are “paradigmatically intentional, characterized by their propositional content.” In other words, experiences are characterized by being, well, *experiential*, or qualitative, while beliefs are characterized by being *about* something. (The trait of a mental state being “about” something is what philosophers call *intentionality*.) But this distinction is already muddled. For it is not just beliefs but perceptual experiences that refer to things—that have intentionality. Chalmers does admit that there is a point where the domains of belief and of perception intersect: namely, in beliefs that are about experiences, as when one believes that one is now having a red experience or that one is experiencing pain. This, however, is another muddle. When I see red or suffer pain this is not a matter of belief. It is not something that is vulnerable to refutation. I cannot mistakenly be in pain.

Chalmers himself recognizes this. He argues that “direct phenomenal beliefs cannot be false”; but in a later chapter he gives reason for thinking that they actually can be. His uncertainty is inevitable given the confusing nature of the phrase “direct phenomenal beliefs.” If we take the view that these beliefs are infallible, then they are, of course, not beliefs—not even “direct” ones, whatever that means—for beliefs may be false. These muddles are so elementary (indeed, they are implicitly acknowledged by Chalmers himself) that they must serve some purpose. This purpose is to advance what he now calls “phenomenal realism”—the view, which we are already familiar with, that phenomenal properties are not conceptually reducible to physical or functional properties.

One way of revealing the autochthonous nature of phenomenal properties is to imagine, as Frank Jackson did in a famous thought experiment known as “Mary’s room,” the case of a super-scientist named Mary who knows everything about the physical properties and functional relations of color but is herself completely colorblind. If, as a result of surgery, she were endowed with color vision, it is obvious that she would have been introduced to something new, since phenomenal experience is different from the material world as revealed to or described in physics. This argument, it would seem, is sufficient to show that subjective experience is not reducible to the kind of objective knowledge that describes the physical world.

---

But Chalmers is not satisfied with this, and has to argue additionally that “pure phenomenal concepts and phenomenal beliefs are conceptually irreducible to the physical and functional because these concepts themselves depend on the constitutive role of experience.” He also refers to a “class of concepts that have phenomenal concepts as constituents.” In other words, he asserts that some concepts and beliefs are irreducible by proxy—by arguing that some of them depend on experience, and experience itself is irreducible. But this is a narrow and unnecessarily complicated way of asserting that we cannot see how such concepts could arise out of the purely material world. Indeed, how do *any* concepts arise out of the inert matter-energy interchanges of physics? Until we are presented with a plausible account of how the *concept* of “matter” arose out of *matter itself*, we should be prepared to argue that there is nothing in matter as described by physics that would suggest it could rise above itself, and enclose that which it has risen above in quotation marks. (It is this simple insight—and not anything about how confusing, difficult, or incomplete is quantum physics—that is one of the great challenges to materialism.)

Once one recognizes the *inherent* irreducibility of concepts, little more needs to be said on the idea of “concepts of consciousness.” Certainly that little more need not occupy the nearly ninety closely printed pages that follow. Chalmers could have spared his readers a grueling trudge, marked from time to time by what he seems to believe are bold assertions—such as that “a wide range of social concepts will turn out to be partly phenomenal,” as if things could be otherwise. Society could hardly be established in the absence of phenomenal consciousness; and social *concepts* seem to exist about as far into the phenomenal realm as they could relative to the “microphysical” truths that he sees as fundamental to physics. Similarly, it hardly seems that Chalmers should introduce as an “intriguing possibility” the notion that “phenomenology could play a crucial role in a subject’s *possessing* a causal or a mathematical concept even though these concepts are conceptually independent of phenomenal concepts.” The proper response would seem to be: you bet.

That Chalmers has to work at rescuing the contents of consciousness from the physical world is the consequence of what we noted at the outset: his allocating too much to the easy problems, ceding to physicalism territory that belongs to irreducible phenomenal consciousness. This is illustrated by his assertion that “many mathematical or philosophical concepts have no obvious tie to phenomenal concepts,” which may be true, but does not mean that those concepts are not ultimately derived from an irreducibly phenomenal consciousness. Imagining that they could be generated

---

by the physical world is the result of misplacing explicitness, so that our concepts are embedded in the material world.

This tendency can go far into madness: Chalmers explores the idea that microphysical processes are constituted by computational processes—in other words, that reality is actually just one giant computation. (This is sometimes called the “it-from-bit” doctrine.) This should not be surprising, since we have already seen how he believes that “information” can be found in electrons and rocks. Even so, it is interesting to note the mirror tendencies: Chalmers has great difficulty acknowledging the phenomenal consciousness of features of the mind that clearly exhibit it, but sees it with great ease in the entire material world, including the parts of it that are not even alive. If one tries to show how things that could not occur without consciousness (things like perception and cognition) actually can be explained as purely part of the material world, then one will have to see the physical world as infused with things like “information”—which in normal, non-technical parlance, requires consciousness and first-person awareness. While *narrowing* conscious awareness to encompass only qualia may seem like an opposite move from *expanding* the presence of consciousness to the entire material world, they are actually both consequences of doing away with the gap between the way we talk about physical events and the way we talk about experience. This failure to maintain the ordinary, intuitive distinctions between the physical and the mental undermines Chalmers’s project at its very heart.

### Representationalism Revisited

It will be recalled that Chalmers rejected representationalism, the theory that a mental state is phenomenally conscious if and only if it *represents* some fact about the world and plays a functional role in transforming the organism’s perceptual inputs into behavioral outputs. Chalmers resists functionalist interpretations of experiential consciousness, arguing that the phenomenal contents of consciousness cannot be reduced to their causal relations, their role in causing behavior. Yet, in the chapter he dedicates to representationalism, Chalmers, as with so many other aspects of the mind-body problem, concedes too much. That chapter begins:

Consciousness and intentionality are perhaps the two central phenomena in the philosophy of mind. Human beings are conscious beings: there is something it is like to be us. Human beings are intentional beings: we represent what is going on in the world. Correspondingly, our specific mental states, such as perceptions and thoughts, often have a phenomenal

character: there is something it is like to be in them. These mental states also often have intentional content: they serve to represent the world.

This passage repays close critique, as the muddle that it evinces is pandemic in today's philosophical discussions of the conscious mind.

Chalmers elsewhere admits that there is a close connection between intentionality and phenomenal consciousness, and notes that there are two schools of thought about this relationship: one holding that consciousness is grounded in intentionality, and the other holding that intentionality is grounded in consciousness. He admits that his sympathies lie with the former, and therefore thinks that there is "significant promise" in the theory of representationalism. It is important to see that these beliefs arise from a crucial error, guided by the misleading metaphor that we find in the passage above: the notion that an intentional relation is one of representation.

Consider a straightforward example from the field of vision. If I look in the mirror, I see a representation of my face. The relationship between my face and its mirror image is purely physical. The image is an *image* by virtue of having physical characteristics in common with my face: it replicates the surface appearance. It is generated by a causal interaction between the surface of my head and the silver of the mirror, mediated by light that has bounced from one to the other. Now consider myself looking at the image in the mirror. My sensory experience *refers* to the image, but is not a replication of it (at least, not in the sense in which a mirror image is a replication of my face). The relationship between my face and its mirror image is a causal relationship, whereas the relationship between the mirror image and my awareness of it is more than a causal relationship. Even if one believes in the neural theory of perception—that neural firings are the cause of perception—there still has to be, in addition to the physical cause, an *intentional* relationship, by which my perception *refers* to the mirror, reaching causally "upstream" from the nerve impulses in the visual cortex to the image that is located in the mirror. But there is no such intentional relationship between my face and its reflection in the mirror; that is a relationship of representation, or potential representation, to be realized in a conscious being, and no more.

We can further underline the difference between representation and intentionality by noting that, whereas the mirror image is a *re*-presentation of my face, my perception of my face is a *presentation*. (My perception of the mirror image, then, is a presentation of a representation.) There is no representation without presentation (or "making present"). And for presentation, a causal relationship is not enough: it requires intentionality. The representational theory of phenomenal consciousness is actually one

---

that Chalmers should eschew, because it brings him too close to the functionalism that would reduce the contents of consciousness to entities in the causal nexus that passes through the brain.

Chalmers is aware of this unwelcome alignment, and tries to wriggle out of it by attempting to preserve the intentionality of representations. This he does first by speaking of “pure” representational properties, which are characterized by “representing a certain intentional content.” But it is not clear how an intentional content could be “represented,” or why it would be necessary to represent an intentional content anyway if it were already sufficient for referring to something.

He tries even harder to rise above the trap of representationalism by suggesting that representational contents *correspond* to states in the world, and that they have “conditions of satisfaction,” meaning that they can either correctly or incorrectly represent those states in the world that they represent. For example, the image of my face in the mirror may display a blemish, but that could be a flaw in the mirror, not in my face; so the image can be verified or falsified as a representation by looking at the thing in the world that it represents. The problem is that mere representations of this sort are simply physical effects of physical causes, and make no claim of satisfaction—unless, of course they lead to expectations *in us* that may or may not be fulfilled. But these expectations are an inferential relationship, and this is what Chalmers implicitly relies on to make his case, when it is just what he needs to be demonstrating.

Chalmers’s entire discussion of “The Representational Character of Experience,” as this chapter is titled, is marred by these confusions. His attempts to distinguish various features of representation, like “pure” representational properties versus “impure” representational properties, the scope of representationalism versus that of phenomenal consciousness, and Russellian versus Fregean contents of perceptual experience (don’t even ask), end up weaving cognitive hairballs. The complexity of this chapter, and the detailed exegesis of these various distinctions, do not liberate it from its original terminological sin. Nor do they deliver this work from embracing the very reductionism Chalmers is setting himself against, which, as he says, understands representation in purely “causal, informational, and teleological terms.” (The aggregation of these three descriptions is notable: “information” is reduced to causal relations, and “teleology” to a sort of causal wiring-in of the organism to the biosphere that ensures its flourishing.)

Ultimately, Chalmers does not persuade that he has made the case for his “nonreductive representationalism,” which says that “phenomenal properties are equivalent to representational properties that cannot be

---

fully characterized in nonphenomenal terms.” This theory is a way of hollowing out the most essential part of experienced consciousness—the experience of it—while acknowledging that it still sticks around as a sort of curious remainder. This, of course, traces back to the fatal first step, running through all of Chalmers’s work, of narrowing the presence of phenomenal consciousness in our minds, so that many things are excluded from the scope of the hard problem. If the main tasks he described at the beginning—in his words, “the integration of information by a cognitive system; the reportability of mental states; the ability of a system to access its own internal states; the focus of attention; the deliberate control of behavior; [and] the difference between wakefulness and sleep”—do not require awareness or internal experience, then it is easy to see how the barrier between phenomenal consciousness and “representation” could be broken down. Most importantly, this allows the causally closed material world to encroach too closely upon consciousness.

The task is actually both much simpler and much more difficult than Chalmers has made it out to be. The quarrel over whether consciousness depends on intentionality or the other way around is empty. Elements of phenomenal consciousness such as perception, conscious beliefs, and other propositional attitudes such as *fearing that* and *hoping that*, are inescapably “about” something or other. The only elements of consciousness that might not be considered intentional are isolated qualia, such as feelings of warmth, tingling, or the color red. Yet even these sensations are dense, layered, and multifaceted; more to the point, these sensations are *located* in, and so refer to, intentional objects: warmth refers to or is *of* a warm arm, or tingling is of a numb leg; the color red refers to a red ball, or to a memory, or perhaps to some inner state of consciousness itself if one is contemplating a Rothko painting. Although his sympathies lie with those for whom intentionality is grounded in consciousness, Chalmers admits that there is something to be said for the opposing camp. But this ought to lead him to the obvious conclusion that each is actually inseparable from the other, and that the opposing views are simply looking at the undivided pair from different angles.

### **Out of the Garden, into *The Matrix***

I can look at my perceptions (so to speak) in two ways: in terms of the objects that I perceive; and in terms of the “what it is like” to perceive the objects. Within the flow of everyday life, this is a distinction without a difference. I can, however, stop to reflect, and dissect my perceptions into components, which are likely to be sensations: I can say that I see a chair,

---

or I can say that I see a patch of brown or a combination of shapes. The latter seems to be closer to what Chalmers calls “phenomenal consciousness,” understood as “what it is like” to have an experience; the former is closer to what he means by the ill-chosen term “representational content.” The relationship between these two is of interest to Chalmers. But a related and more interesting question is opened up by considering what we might call the layeredness of consciousness: the fact that, all at once, we have the experience of sensations; the experience of our body as the bearer of those sensations; the experience of ourselves as bearers of those sensations; the experience of embodiment; the experience of the environment of our bodies; and the experience of what I have called (in *The Knowing Animal* [Edinburgh, 2005]) “propositional awareness,” our sense that certain states of affairs are the case—a sense that encompasses memories, awareness that such and such is present before us in the world, facts, and so on.

Chalmers’s starting point seems to be an attempt to assimilate awareness, phenomenal consciousness, experience, and the rest into the single thing that Thomas Nagel calls the “*something it is like* to be a conscious organism”; and while consciousness is indeed unified, this way of looking at it unhelpfully collapses all of the many layers of consciousness, gathering up contents as disparate as having a tingling sensation, entertaining a belief, and remembering a historical fact, but without carefully preserving their distinctiveness. Experiencing a toothache, seeing a horse, worrying about an exam in a fortnight, and remembering the date of the Battle of Hastings are all tinged with phenomenal consciousness, but none of them is fully captured by the “what it is like to be” formula. It is this that prevents Chalmers’s conceptual inquiry from getting beyond detailed and circular discussions of, say, “the interface between consciousness and intentionality.” And his discussion of the contents of consciousness is equally impoverished.

The mist clears briefly in the opening to a dense, seventy-page chapter enticingly titled “Perception and the Fall from Eden.” In the Garden of Eden, Chalmers says (tongue in cheek), we had unmediated contact with the world. Then there was the Fall. We discovered that the appearance of things varied even when those things in themselves did not change: “we ate from the Tree of Illusion” and realized that there is a contingent connection between what is in the world and how we experience it. And then “we ate from the Tree of Science” and discovered that our experiences are mediated in a more complex way by physical processes.

But this foundational tale leads back into a discussion of the relationship between the phenomenal character of perceptual experience and its representational content, which seems like an endless circling around

---

something that never comes clearly into view. We are repeatedly told things we probably do not need to be told, such as that “when a certain book appears red to us, there is a quite reasonable sense in which the experience will be satisfied [if and only if] the book in question is red at the relevant location.” And even if his notion of the “primitive” were clear, Chalmers’s conclusion that “consciousness may consist in the phenomenal representation of certain primitive properties” would be a disappointing outcome, especially as the end of an interminable investigation of (again) Fregean versus Russellian understandings of meaning.

If, as Chalmers seems to be saying, it is “Edenic” content in which we had unmediated contact with the world, then this is extremely puzzling. Even someone as hostile as myself to materialist accounts of consciousness cannot subscribe to the notion that we are “directly acquainted with objects in the world,” that objects are “presented to us without causal mediation,” or that Edenic content is “the most fundamental sort of content of an experience.” After all, we always see objects from a point of view that is ours, not theirs; if we manage to escape this limitation it is only by withdrawing from experience to measurement, to a view from nowhere in which the object no longer has a phenomenal appearance (and even this is only appearance at a great remove). In short, the idea of an Edenic appearance is a contradiction in terms.

But in the background of this discussion is the ever-present, crucial question as to whether perceptions really give us access to an outside world. Is it possible we might be systematically deceived as to the seeming independent existence of the things we are conscious of? Are they in fact internal to consciousness, so that the world of objects out there is an illusion? Chalmers examines this question through the famous thought experiment dramatized in the cult film *The Matrix*: a brain is floating in a vat of nutrients and is stimulated, with the same sort of electrical inputs that a normally embodied brain receives, by a giant computer simulation of the world. The brain is entirely and successfully deceived. How do I know that my situation is any different? Hilary Putnam, who first introduced this thought experiment, argued (decisively, in my opinion) that the hypothesis that I am and always have been a brain in a vat can be ruled out in principle, because “brain” refers to (real) objects in my perceived world, as do “vat,” “nutrients,” and the “scientist” who set the whole thing up. The thought experiment presupposes the existence of the very things that it is supposed to call into question.

This would seem to leave the thought experiment for dead, but Chalmers squeezes forty pages of argument out of its corpse. He claims

---



that it points to the possibility of a non-skeptical, positive three-part “metaphysical hypothesis.” This hypothesis proposes: first, that micro-physical processes throughout space-time are constituted by underlying *computational* processes (a daft idea that we have already touched on); second, that physical space-time and its contents were created by beings outside physical space-time (as many religions hold); and third, that our minds are outside physical space-time but interact with it (as Descartes argued). If any or all of these possibilities turn out to be true, then that would undermine the real existence of chairs, tables, and bodies, but would mean that their fundamental constituents are bits of information, and originate outside of space and time.

Chalmers calls these ideas “a creation myth for the information age.” Quite so—but he offers neither arguments nor evidence for their truth. And the very use of the word “simulation” to make the case undermines it: if everything is a simulation, then nothing is, not least because there would be no first-order reality to simulate. The philosopher Gilbert Ryle argued against drawing from the fact that perception *often* deceives us the conclusion that it *always* deceives us, saying, “there can be false coins only where there are coins made of the proper materials by the proper authorities.” In other words, we could have no concept of illusion if all of our perceptions were illusions; nor could we have a concept of illusion without some of our perceptions being correct. The same rejoinder applies to the metaphysical hypothesis.

More to the point, abstract computation lacks the actual content that is revealed (truly or not) in perception. A computer simulation, which is merely a mathematical shell of the relationship between physical processes, could no more deliver those physical processes than it could, without a subscription to the defunct computational theory of mind, deliver conscious contents. Our simulations of the weather do not create rain inside our computers. Chalmers’s toying with what we might call “pansimulationism” brings him close to the panpsychism he entertained in *The Conscious Mind*, and this, when considered in conjunction with his sympathy for dualism, comes close to justifying some of the criticism that volume received.

### **The Hard Problem Will Not Be Easy**

How much further does this new book, of nearly six hundred closely printed, closely argued pages, take us beyond where Chalmers left the story in *The Conscious Mind* nearly fifteen years ago? Alas, it seems, not much further. The work he—and we, his readers—have to do is not justified

---

by the conclusions he reaches, which are often trivial and usually provisional. The endless discussions of responses-to-objections-to-responses-to-responses give one the sense of an interminable in-house and inward-looking philosophical argument, in which the voices often communicate in shorthand. While it is entirely proper that any conclusions should be accompanied by a full presentation of the “working out” that led up to them, *The Character of Consciousness* seems less a work in progress than a work without progress. We reach the end without a clear idea of what has been achieved and what remains to be done. This is in part because the book is essentially a collection of papers Chalmers has published over the last decade. As a result, there is much repetition—admittedly this sometimes gives the reader a much-needed breather—and all of the to-ing and fro-ing undermines any sense of coherent direction. This is all the more reason that Chalmers ought to have offered a final summary setting out what had been gained, and the direction for further inquiries.

Most disappointingly, the excitement at the mystery of consciousness that glowed through Chalmers’s first book is virtually extinguished. In part this is due to the fundamental error that we have seen pervade the book, of drawing the distinction between the territory of the “hard” problem and that of the “easy” problem in such a way as to hand over too much of consciousness to functional inputs and outputs handled by “cognitive mechanisms.” Consciousness is rather emptied, and the sense of its being reduced to an abstraction is exacerbated by the scores of pages devoted to discussions of technicalities.

This is a great pity. For the seemingly inescapable failure of neural or materialist accounts of consciousness opens up a world of intellectual possibility. As Jerry Fodor has suggested about our attempts to understand consciousness, “The revisions of our concepts and theories that imagining a solution will eventually require are likely to be very deep and very unsettling.... There is hardly anything we may not have to cut loose from before the hard problem is through with us.” Indeed, this seems almost certain, given how inadequate we already know our scientific orthodoxies to be at accounting for the truths of consciousness. Just as rethinking the nature of light transformed our understanding of the physical world, shattering seemingly secure theories of physics to give rise to relativity theory and quantum mechanics, when we are finally able to account for the unfathomable depths of our own minds, it is sure to have profound and transformative consequences for our understanding of what kind of world we live in, and what manner of being we are.