

©1996 The Institute of Mind and Behavior, Inc.  
The Journal of Mind and Behavior  
Autumn 1996, Volume 17, Number 4  
Pages 391-398  
ISSN 0271-0137

**The Conscious Mind: In Search of a Fundamental Theory.** David Chalmers.  
New York: Oxford University Press, 1996, 432 pages, \$29.95 hard.

*Reviewed by Valerie Gray Hardcastle, Virginia Polytechnic Institute and State University*

---

David Chalmers' book *The Conscious Mind: In Search of a Fundamental Theory* is well-written, though a bit repetitious. He follows the current major arguments for why materialist theories of consciousness can't work and then advances his own dualistic theory of consciousness based on Shannon information partitions. There is much — probably too much — territory covered in this book, and in this review I hope to present a fair summary of what Chalmers believes and to offer some reasons why his approach is not the best one. (I mention this aim because many of the reviews of this book focus only on his recounting of the already well-known thought experiments for why materialism is odd; here I shall focus on other things.)

Chalmers starts several chapters as philosophically technical and suggests that the general reader can skip or skim these sections without loss. Though I agree with his assessment that the chapters are technical, I think it unfortunate that some may not read these portions of his book, for it is here that Chalmers makes what I think are his deepest mistakes. To appreciate this claim, though, we need to start with what he does believe.

In chapters six and nine, Chalmers explores in detail the "remarkable coherence between conscious experience and cognitive structure" (p. 218). He readily acknowledges that there are neuronal and information processing correlates of consciousness, much as all materialists do. Indeed, he takes these parallels as being fundamental in developing a good scientific theory of consciousness. I do disagree with which processes he picks out as being correlated with conscious experience (the contents of the mind available for global control), but I take this to be a small matter in his overall framework. Nothing he says or believes hangs on this connection and I would think that Chalmers would be perfectly happy to go with whatever scientists ultimately discover as the neural or psychological correlates of consciousness.

For many, discovering such correlations are what identity hypotheses are made of. We note that whenever we have water, we also find H<sub>2</sub>O and eventually we

---

I would like to thank David Chalmers, Owen Flanagan, Gary Hardcastle, and Gregg Rosenberg for their discussions of David's book and their comments on an earlier draft of this review. Requests for reprints should be sent to Valerie Gray Hardcastle, Ph.D., Department of Philosophy, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061-0126.

conclude that water must be  $H_2O$ . We use “remarkable coherence” as the best evidence around for identity claims. That is just how science functions, for the most part. So why isn’t Chalmers a materialist? Though he does have a rather naive view of the complexities of science and scientific explanation — he systematically ignores the pragmatic dimensions of explanations, for example — he doesn’t gloss over this aspect of scientific practice. Indeed, he makes heavy use of it in the latter half of his book. The answer turns on the details in starred portions of chapters one and two. Let me sketch the relevant bits.

There are two types of meaning words can have. The *primary intension* refers to how we common-sensically think of a term. So the primary intension of “water” is “that clear potable stuff” or something like that. The *secondary intension* really comes from science. We investigate that clear potable stuff in this universe and discover that all instances of it coincide with instances of the structure  $H_2O$ . So we conclude, after some experimentation, that water is  $H_2O$ . Notice that water did not have to be  $H_2O$ , or so says Chalmers and a whole host of other philosophers. It could have been XYZ, had the universe been structured differently. The identity between water and  $H_2O$  is purely contingent; it depends on how the universe actually is. This identity contrasts with the primary intension of water, which, according to Chalmers and others, is true *regardless* of how the universe actually is. The term “water” refers to clear potable stuff, independently of the underlying structure of the stuff. As Chalmers puts it, this is a “conceptual” connection. These identities hold necessarily, for they turn on how our language is arranged, not on how the world is. As long as we speak the same language, then the primary intension of water will be the same and we can know the primary intension without investigating the world at all. All we have to do is know the concepts.

So too for the concept “consciousness.” Our primary intension of consciousness refers to our subjective experiences. After we investigate the world, we might discover that these experiences are identified with or as something, but this would only be a contingent identity, a secondary intension of the term. Had our world turned out differently, then consciousness might have turned out to be identified with something else, just as water might have turned out to be something else too. (Chalmers apparently doesn’t agree with this line of reasoning. He maintains that “with consciousness, the primary and secondary intensions coincide” [p. 133] and they both are tied to having a phenomenal feel. But, given that he argues for a specific dualistic conception of consciousness that goes far beyond even his original intuitions, he can’t be entirely serious here. The secondary intension could at least be richer than the primary.)

Chalmers leans hard on these claims, asserting that if we want to take consciousness “seriously” (p. xii), then we have to explain *why* our brains are the things that have conscious experiences. That is, we have to explain *why* the secondary intension means what it does. Here, I believe, is where most materialists (myself included) part company with Chalmers. Standard materialist lines in response to Chalmers’ assertion include: (1) We don’t make that sort of demand on our other identity claims. We don’t have to explain *why* it is that water is  $H_2O$ . It just is for our universe. (2) Science is just in the business of articulating secondary intensions. Once we find the identities, then our job is over. Seeking some resonance with a priori conceptual intuitions is a silly philosopher’s game. (3) Chalmers’ intuitions about the primary intension of consciousness is flawed and certainly should not be used to guide scientific inquiry or theory-building.

Though I have great sympathies with all those lines, I also think that Chalmers' arguments fail on their own terms, for two reasons. First, as Chalmers himself discusses, in the early 1950's Willard van Orman Quine argued that any distinction between conceptual connections and empirical ones is artificial and misleading. Quine held that, given enough and the right sort of pressure from the environment, we would be willing to alter *any* identity statement. We can do this because "our statements about the external world face the tribunal of sense experience not individually but only as a corporate body" (1951, p. 41). By making sufficient changes elsewhere in our web of beliefs, then we can either hang onto or rid ourselves of any belief statements, including Chalmers' conceptual identities.

I haven't the space here to defend Quine's claims. Suffice to say there is an extremely rich literature covering his arguments and that it at least needs to be taken seriously when running the sort of line Chalmers wants to. Though he does address Quine's worries, he does not do so successfully.

In defending his claim that primary intensions are completely a priori against the meaning holists who maintain that no meaning can be set absent experience, Chalmers argues that his particular conceptual claims are in fact beyond the tribunal of experience, for they are *supervenience conditionals* and not identity statements. A *supervenes* on B, where A and B are facts, if B determines A. So, for example, the law of Darwinian natural selection supervenes on the laws of physics, because the physical laws (plus initial conditions) determine how biology unfolds. Materialists generally hold that (descriptions of) the mind supervenes on (descriptions of) the body, because the body determines the mind. (These sorts of claims are weaker than straight identity claims because, for example, they would allow different bodies to instantiate structurally identical minds. A computer and Valerie Hardcastle could process information in the same way, even though we are made out of radically different stuff. So aspects of neither the computer nor Valerie Hardcastle would be *identical* with the information processing since nothing about the computer is identical to the human, or so one line of thinking goes.) The supervenience claim that Chalmers relies on is something like: "If the low-level physical facts about the universe are as they are, then consciousness exists." He believes that this sort of statement is immune to Quinean considerations. It is worth quoting Chalmers' argument in full:

The facts specified in the antecedent of this conditional effectively include all relevant empirical factors. Empirical evidence could show us that the antecedent of the conditional is false, but not that the conditional is false. In the extreme case, we can ensure that the antecedent gives a full specification of the low-level facts about the world. The very comprehensiveness of the antecedent ensures that empirical evidence is irrelevant to the conditional's truth value. (p. 55)

This is simply wrong-headed. Quine's point is that there might be circumstances in which we would want to alter how we think of consciousness. This claim is independent of the truth value of the conditional. Suppose Daniel Dennett became extremely fashionable and his arguments that there are no such things as qualia won the day. Over time, we alter our everyday conception of consciousness to mean reportability. Suppose David Chalmers becomes king and decrees that consciousness must refer to a nonphysical aspect of Shannon information. In response, we alter our everyday conception of consciousness to include dualism. Neither possibility is very likely, of course. My only point is that the proper focus for Quine's denial

of the immutability of primary intensions is with the consequent in Chalmers's conditional, and Chalmers gives us no reason throughout his text for why his or my or anyone else's intuitive notion of consciousness might not change over time. Indeed, he acknowledges that others may not share his intuitive notions to begin with. All they have to do is convince him that their view of consciousness is right and Quine wins the day. The primary intension changes and how Chalmers is supposed to run all his thought experiments to show that there is no necessary connection between consciousness and matter is now unclear. As long as the primary intension is mutable, then we should not put much weight on arguments that depend upon them. Certainly, they should not be allowed to determine ontological "facts." Better to see that our concepts evolve as our understanding of the world does.

Second, Chalmers' analysis of logical necessity entailed by primary intensions is wrong. He suggests that there are two ways in which A can supervene on B: naturally or logically. A natural or nomic supervenience refers to items only in this universe, constrained by our natural laws. A supervenes naturally on B if, according to the laws and conditions of this universe, B fully determines A. A logical supervenience refers to any logically possible universe, where logically possible roughly means conceivable. A supervenes logically on B, if regardless of the laws of the imagined world, B fully determines A. Chalmers does believe that consciousness supervenes naturally on the physical, such that there is no way for consciousness not to arise, given our physical laws and our universe's initial conditions. Another way of putting this point is that he accepts that an explanation of consciousness based on its secondary intension will include physical beings and their interactions. However, he denies that it supervenes logically. That is, he can imagine worlds in which there is the same physical stuff as this universe, but no consciousness. These are all the Zombie worlds that philosophers are so fond of discussing. The primary intension of consciousness (as Chalmers currently understands it, at any rate) is not tied to any particular physical interaction or stuff.

But this isn't the whole story, for there is a problem with the logical possibility of "a world physically identical to ours, but with additional nonphysical stuff that is not present in our own world: angels, ectoplasm, and ghosts, for example" (p. 39). If these angels follow biological laws, say, and differentially reproduce and evolve, then biology might not logically supervene on the physical. But, as Chalmers says, "we certainly want to say that biological properties are [logically] supervenient on physical properties . . . . Intuitively, it seems undesirable for the mere logical possibility of the angel world to stand in the way of the determination of biological properties by physical properties in our own world" (p. 39). So we need to restrict our notion of logical supervenience a bit.

Chalmers offers two restrictions. First, we "turn supervenience into a thesis about *our* world (or more generally, about particular worlds)" [p. 39]. So A-facts logically supervene on B-facts if in any possible world with A-facts, at least the B-facts will be true. Any additional B\*-facts then (the existence of angels, say) will not count against the supervenience relation. Second, since this restriction doesn't help with the supervenience of certain general facts in our world — that there is no such thing as ectoplasm, for example — Chalmers disallows negative claims (this would also include universal statements such as all kangaroos are mammals). "Supervenience theses should apply only to *positive* facts and properties, those that cannot be negated simply by enlarging a world" (p. 40). And here, I believe, Chalmers ties himself up in knots.

What counts as an A-fact? Chalmers has in mind something like all the low-level facts about the physical world, a description of our universe as points in Hilbert space, say. He wants these facts to stand apart from any contingent nomic connections between the A-facts and higher level descriptions. He thinks the existence of  $H_2O$  an A-fact about our universe and the nomic connection between  $H_2O$  and surface tension as a contingent law in our universe. But notice that, since our universe is finite, I could redescribe all the laws in our universe as particular statements about particular things. Conceivably, we could articulate all the “contingent nomic connections” as local relational facts about the lower-level physical stuff in our universe, as A-facts, in other words. If this is the case, then there is no real distinction between logical and natural supervenience, as Chalmers describes it. Anytime he wants to hold *all* the A-facts of our universe constant, he is also holding constant all the instances of the “contingent” facts too. Consequently, if he thinks that consciousness supervenes naturally, then it also must supervene logically, and he can’t imagine worlds just like ours, only lacking subjective experience. (Chalmers actually runs his discussion in terms of properties instead of facts, but the end result will be the same: I can redescribe all the nomic connections between physical stuff and consciousness in terms of particular relational properties of the lower level stuff.)

To his credit, Chalmers recognizes that architecture, astronomy, behavior, biology, chemistry, economics, meteorology, and sociology — to name a few sciences — all supervene logically on the A-facts, according to his definitions (p. 73). He just thinks that consciousness is plain different. But it can’t be, if it supervenes naturally in our world, for the relation between it and physics is already fixed, just as the relations between biological creatures or economies or meteors and physics are already fixed, no matter what he thinks about them a priori.

But maybe these are picky philosopher’s points, better relegated to the starred sections of his book. Suppose Chalmers could defend a conception of consciousness that would be untouchable by empirical considerations and that he could come up with a notion of supervenience that sustains a distinction between the logical and the natural. How does one go from acknowledging that the primary intension of consciousness does not require physical stuff and that we can truly conceive of worlds like ours only lacking experience to being a dualist? Chalmers sees this as a natural move. I am not so sanguine.

If Chalmers’ analysis is right, then the lawlike connection we see between consciousness and brains is a further fact about the world, not determined by the physical facts alone. He suggests we just have to accept this further fact as something brute. There is no explanation for it; it just is the way the universe is, similar to how we accept the connection between water and  $H_2O$ . (I note for the record that his position belies his original task, taking consciousness seriously by explaining why it is that certain things are conscious in our universe. Accepting some fact as brute certainly isn’t explaining it. Does this mean Chalmers doesn’t take consciousness seriously after all?)

So what is this further fact we are accepting as brute? As I see it, there are two ways we could go. On the one hand, we could say that the further fact we have to accept is that consciousness just is what it is nomically correlated with, much in the same way we say water just is what it is correlated with. So consciousness would be some physical thing in the universe after all. This might be a “merely” contingent identity, but it is just something we have to accept as a brute fact about the world, given the way science turns out. On the other hand, we could say that, contrary to

the established scientific method for discerning identities, here all we get are correlations, nothing more. Consciousness is nonmaterial and we have to accept the unusual connection between certain physical configurations and this nonphysical stuff as brute. Materialists opt for the first line of reasoning. Chalmers for the second. For the life of me, I can't see why, except that it saves his intuition that consciousness isn't material. (On the other hand, Chalmers himself points out that "finding a conclusion counterintuitive or repugnant is not *sufficient* reason to reject the conclusion" [p. 159].) But the mere fact that we are going to have to accept an additional fact about the relation between consciousness and the world as brute does not entail dualism, despite Chalmers' assertion to the contrary.

Actually, I think Chalmers slips on his use of the notion "physical" here. In the first chapters of the book, he explicitly says that it refers to physics and its positings, not all the stuff in the universe. So quarks and bosons are physical, in his sense, and tables and chairs aren't. But by the time he gets to chapter four, he seems to go back to using physical in its broader sense. He writes: "the presence of consciousness is an *extra* fact in our world, not guaranteed by the physical facts alone" (p. 123) and from that he concludes that "materialism is false" (p. 123). But he could only conclude that if he means physical here in a very broad sense, in which case both statements amount to the same thing. *Prima facie* there is no reason that the extra fact about the universe not entailed by physics couldn't be physical (in the broad sense) as well. (We can see this point demonstrated in quantum mechanics. There are facts in this world that the physical facts don't entail, because their appearance is probabilistic. Nonetheless, they still refer to physical things in our world.)

Indeed, the dualism Chalmers feels forced into defending gets so bizarre (and unsubstantiated), that I do wonder why he doesn't go back and re-examine his choice at the fork. He ends up advocating a kind of panpsychism in which he sees consciousness everywhere and as solving the deep mysteries of quantum mechanics. His model for understanding phenomenal experience draws from Shannon information theory. He notices that our cognitive states and our phenomenal states both map well onto the difference relations in an information space, so he supposes that information might be the glue that ties our psychological processing to our phenomenal experiences. He also notices that Shannon information is everywhere, so, he concludes, consciousness must be everywhere too. Finally, he notes that our basic physical theories are, at bottom, relational, "a pure causal flux, with no further properties for the causation to relate" (p. 153). Since phenomenal experience is only one thing we have any "direct familiarity" with that is intrinsic and nonrelational (p. 153) [of course, it is the only thing we have any direct familiarity with, period], this must be doing the relational work for physics. This feeds into a highly unlikely interpretation of quantum mechanics in the final chapter, which Chalmers himself confesses that he cannot "wholeheartedly believe," for it is simply "too strange" (p. 357).

Needless to say, I find this sort of argumentative strategy dubious. Here is an analogy for it: I notice that bubble-gum has mass. And there are properties of mass that map well onto properties of bubble-gum (divisibility, location, shape, duration, etc.). I also notice that chicle resin (which is what gum used to be made out of, at any rate) has the same sort of properties. Hence, I conclude that mass must be the glue that ties chicle resin to bubble-gum. But then I notice that lots of things in the world have mass. So, everything that has mass must be bubble-gum! Better, I would think, to use the discovered correlates as grist for the identity mill (chicle resin is

gum under certain circumstances; phenomenal experiences are psychological processes under certain circumstances), than to take the unifying descriptions as some underlying third cause (or non-cause, if you are in the phenomenal realm). At the least though, Chalmers' panpsychism does fit with the work consciousness must do to save physics.

Finally, Chalmers' view cannot escape the old problems with garden-variety versions of dualism, which he tries to address in his chapter five, "The Paradox of Phenomenal Judgments." If dualism is true, then my claim that I am seeing red now has no obvious connection to my conscious experience, for we can account for that claim solely in terms of a photon triplet impinging on my retina, various brain processes being executed in my head, and then some muscle contractions in my face along with the expulsion of carbon dioxide. My being aware of red has nothing to do with it. As he writes: "Consciousness is conceptually independent of what goes into the explanation of our claims and judgments about consciousness" (p. 179). Here Chalmers bites the bullet, claiming that we need to "separate what is *merely counterintuitive* from what threatens the viability of a nonreductive view of consciousness" (p. 181, italics mine). [Why that reasoning isn't used in the first half of the book to disallow his own use of an intuitive notion of consciousness is unclear.] It is indeed counterintuitive that my conscious states have nothing to do with my judgments about my conscious states, but that "does not matter" (p. 188), for being odd is not enough to defeat a theory. (This, of course, is the line the materialists use to defend their own theories, but I don't suppose they have squatter's rights. However, it does seem unfortunate that Chalmers rejects that line of argument when arguing for his own approach but then relies on it in defending his approach from criticism.)

Indeed, our knowledge of conscious states is quite different than any other sort of knowledge. As Chalmers argues, it isn't causal or reliabilist, but relies on "another sort of relationship entirely" (p. 193). We have access to our conscious experiences "directly" (p. 196), and our "*having* the experiences" justifies our beliefs about them (p. 196). This is not the same as *knowing* the experience automatically (p. 197), for we can be mistaken in our beliefs about our conscious states. How having this sort of fallible yet direct access to something differs from reliabilism escapes me. Chalmers just asserts that "there is something special about our knowledge of experience" as a point of fact (p. 197), which isn't going to sway a skeptic. He concludes the chapter by hoping that he has "said enough to make clear that [his] . . . view provides a natural framework for making sense of these issues" (p. 209), but I confess that without more details of what "having direct access to" amounts to, I see little distinction between Chalmers' use of "natural" and my use of "miracle."

Chalmers wants a "naturalistic dualism" (p. 128), a dualism that is perfectly amenable to scientific investigation and study. However, he himself notes that "nothing [in his view] contradicts anything in physical theory" (p. 128), where physical theory just is a completed physics (or maybe all of science, depending on how he is using the term). Though I recognize we have moved beyond the simple falsifiability criterion for science, theories still have to make some sort of empirical contact with the world. His dualism would be much more respectable if it did contradict something publically observable, so we could test its claims. At bottom, I can't see "naturalistic dualism" as anything but oxymoronic.

Nevertheless, Chalmers gives it a good shot. His is the best, clearest, and most thorough defense of dualism I have seen. Along the way, he discusses many related and complicated issues in philosophy of language, metaphysics, and artificial intel-

ligence. (His discussion of quantum mechanics should have been omitted, or expanded into a book in its own right.) I would recommend his book to anyone who wants to get a handle on the state of the art in consciousness studies. I just would not recommend it as gospel.

### Reference

Quine, W.v.O. (1951) Two dogmas of empiricism. *Philosophical Review*, 60, 20–43.