Dialogues in Philosophy, Mental and Neuro Sciences

Crossing Dialogues Association

ORIGINAL ARTICLE

I can't get no (epistemic) satisfaction: Why the hard problem of consciousness entails a hard problem of explanation

BRIAN D. EARP

Department of Experimental Psychology, University of Oxford (UK)

Daniel Dennett (1996) has disputed David Chalmers' (1995) assertion that there is a "hard problem of consciousness" worth solving in the philosophy of mind. In this paper I defend Chalmers against Dennett on this point: I argue that there is a hard problem of consciousness, that it is distinct in kind from the so-called easy problems, and that it is vital for the sake of honest and productive research in the cognitive sciences to be clear about the difference. But I have my own rebuke for Chalmers on the point of explanation. Chalmers (1995, 1996) proposes to "solve" the hard problem of consciousness by positing qualia as fundamental features of the universe, alongside such ontological basics as mass and space-time. But this is an inadequate solution: to posit, I will urge, is not to explain. To bolster this view, I borrow from an account of explanation by which it must provide "epistemic satisfaction" to be considered successful (Rowlands, 2001; Campbell, 2009), and show that Chalmers' proposal fails on this account. I conclude that research in the science of consciousness cannot move forward without greater conceptual clarity in the field.

Keywords: hard problem, consciousness, explanation, Dennett, Chalmers

DIAL PHIL MENT NEURO SCI 2012; 5(1): 14-20

INTRODUCTION

David Chalmers (1995) suggests that there are multiple "easy" problems of consciousness, but only one truly "hard" problem, namely that of explaining what philosophers term qualia, or the felt aspects of subjective experience. Crucially, Chalmers argues that these disparate problem-types require explanations that are different not just in degree of difficulty, or in precision of detail, but *in kind*. Specifically, the reductive[a], functional[b], and physical-mechanical explanations offered by cognitive science, he claims, whilst adequate (in principle) to solve the easy set of problems, are by their very nature impotent to address the hard one. In a rebuke, Daniel Dennett (1996) describes this categorical divvying-up between easy and hard as a "major misdirector of attention, an illusion generator" and hence not "a useful contribution to research" (p.4). Instead, Dennett thinks that any explanation adequate to solve the easy problems would leave exactly nothing further to be explained: there *is* no "hard problem" above and beyond the easy ones.

Given these conflicting views, I will join the debate, in this paper, with two arguments of my own. First, I will defend Chalmers against Dennett on the point of his main charge: I think that there *is* a hard problem of consciousness; I think it is distinct in kind from the easy problems; and I think that it is vital for the sake of honest and productive research in the cognitive sciences to be clear about the difference. But I have my own rebuke for Chalmers on the point of explanation. Chalmers, in this same 1995 essay, and in extensive later work (including Chalmers, 1996), proposes to "solve" the hard problem of consciousness by *positing* qualia as fundamental features of the universe, alongside such ontological basics as mass and space-time. But this is no solution. The hard problem of consciousness is a problem of *explanation*, and to posit is not to explain. That is my second argument. To bolster this view, I borrow from an account of explanation by which it must provide "epistemic satisfaction" to be considered successful (Rowlands, 2001; Campbell, 2009), and show that Chalmers' proposal fails on this account. I conclude that research in the science of con-

sciousness cannot move forward without greater conceptual clarity in the field.

DISCUSSION

The hard problem of consciousness

The hard problem of consciousness has been stated in a hundred different ways[c], including in one version at the top of this paper; but here it is again to be clear. It is the problem of *explaining* how it could be that a lump of mere matter, even so intricate and wonderful a lump as a brain, could somehow give rise to, account for, or be otherwise connected or associated with the full richness of subjective, felt experience—that first personal "what-it-is-like-ness" (apologies to Nagel, 1974) of being a being. In order to understand why consciousness in *this* sense is so hard to explain, it may be useful to contrast it with some other senses of consciousness, specifically those whose explanations fall, according to Chalmers, at least in principle within the purview of cognitive science and within the bounds of a physicalist theory of the world.

To wit, the hard problem of consciousness is *not* any of the following. It is *not* the problem of explaining how the mind is capable of discriminating, categorizing, and reacting to environmental stimuli; it is *not* the problem of giving an account of the verbal reportability of mental states; it is *not* the problem of developing a model of the deliberate control of behavior, the focus of attention, or the ability of a system to access its own internal states; and it is decidedly *not* the problem of discriminating between wakefulness and sleep (Chalmers, 1995, p. 200).

All of *those* problems, Chalmers contends, are the problems of a cognitive process in search of a mechanism. In other words, each phenomenon of this type is "functionally definable" (p. 202; see also Howell and Alter, 2009), that is, definable in terms its physical, mechanical components and their interactions which, together, are capable of carrying out the phenomenon's signature function. Hence to explain them, you need only to describe a cognitive, neural, or computational mechanism that is up to the job[d]. Contrastingly, what makes the hard problem so hard—and actually different in kind—is just this: there *is* no mechanical or reductive explanation of a cognitive process, no matter how detailed or complete, that can answer the further niggling question of *why* said process should be accompanied by felt experience. Why doesn't it all just happen "in the dark" as it might in a well-designed robot or in a civilized and well-spoken zombie? What is the difference between us and them? As Chalmers (1995) puts it: "Why should physical processing give rise to a rich inner life at all? It seems objectively unreasonable that it should, and yet it does" (p. 201). That is the problem.

Dennett's rebuke

Dennett (1996) agrees that this is a problem; he just doesn't think that it is different *in kind* from the "easier" problems whose explanations lie within the grasp of cognitive science. His main strategy for making this point is to appeal to history, to a time when earlier philosophers posited a strict line between certain types of phenomena, along with their accompanying explanations, and turned out to be wrong.

"Imagine some vitalist," he writes, referring to a much-maligned (in contemporary thinking) class of non-mechanist thinkers, "who says to the molecular biologists [that] the easy problems of life include those of explaining [things like] reproduction, development, [and] growth," but that the really hard problem is "life itself". Dennett's deliberately dopey vitalist continues with this: "We can imagine something that was capable of reproduction, development [and so on] but that wasn't, you know, *alive*" (p. 4). His point, of course, is that all it *means* to be alive is to be capable of some reasonable conjunct of those subsidiary functions: there *is* no "hard problem" of life above and beyond reproduction, growth, and all the rest. With the smug assurance of hindsight, we are tempted to chuckle at Dennett's example, because we all know who turned out to be correct: not the vitalists with their postulated life-force or *élan vital*, but the hard-nosed, reductive molecular biologists. As Koch (2009) writes, making the same point in a slightly different way: "philosophers deal in belief systems and personal opinion, not in natural laws and facts. They ask interesting questions and pose challenging dilemmas, but they have an unimpressive historical record of prognostication" (p. 392). Give us time, they seem to be

saying: science will figure it all out.

But is this allusion to the "vitalist" gaffe fair, or even helpful? My view is no — not at all; it completely misses the point. In the case of consciousness, qualia are not something anyone posits — like a life force — in order to explain some *other* phenomenon, such as a given cognitive capacity (or in the vitalist analogy, something like metabolism); rather they are something *known* to exist through direct experience, indeed they are *defined* as experience, and *they* are the explananda in need of an explanans.

Chalmers (1995), for his part, anticipates the 'vitalist' objection anyway, and deals with it head on: "If someone says," he writes, "[that they] can see that you have explained [some aspect of life such as] how DNA stores and transmits hereditary information from one generation to the next, but [that] you have not explained how it is a *gene*, [then] they are making a conceptual mistake. All it *means* to be a gene is to be an entity that performs the relevant storage and transmission function" (p. 203, emphasis added). In contrast, someone who acknowledges that you have explained some cognitive function, but thinks that you have failed to explain why there is "something it is like" to undergo, perform, or experience that cognitive function, is *not* making a conceptual mistake. Instead, that person is asking a "nontrivial further question" (p. 203).

In short, Dennett's analogy fails on exactly the point of contention. He does not seem to grasp the difference between a clear conceptual error — as committed by the vitalists, or as illustrated by Chalmers in his example with the gene — and the "nontrivial further question" entailed by the existence of qualia. He is wrong, therefore, to dismiss the hard problem as a non-problem, or to class it as being no different from the "easier" problems of explaining cognition. But he might still be right that Chalmers' distinction between easy and hard is a "misdirector of attention" or otherwise non-useful for research (Dennett, 1996, p. 4). I think that this is not the case, however, as I explain in the next section.

Chalmers' distinction is useful for research

If Chalmers' easy vs. hard account were good for nothing else (and it is good for more than this), at the very least it would serve as a reminder for cognitive scientists and other researchers in the field to be clear about what they mean by "consciousness" when they publish their papers. This is not inconsequential. As Kuijsten (2009) reports, "at conferences on consciousness, it is often the case that no two speakers seem to be talking about the same subject" (p. 2). Or as Chalmers (1995) himself puts it, given the current state of affairs, "those who talk about 'consciousness' are frequently talking past each other" (p. 202). Indeed, Vimal (2009) identifies at least *forty* distinct meanings of "consciousness" in contemporary scholarly use alone.

That so many uses of the word are in circulation may be due to the fact that consciousness is, as Sloman (2009, 2010) puts it, a "polymorphic" concept (in the sense of Ryle, 1951). This means that its meaning can be determined only with reference to the context in which it is used or the propositions with which it is paired. Hence, Sloman urges, researchers should not give 'unitary' explanations "of how 'it' evolved, or how the brain produces 'it', nor [should they discuss] a time at which 'it' first exists in a foetus, nor [provide] a machine model of 'it'" (2010, p. 119). This is a crucial message. The take-home lesson is that scientific progress on "consciousness" cannot move forward a single inch until researchers begin to clarify, in every instance, exactly what they mean when they use the term.

Of course, if simple conceptual clarity were the end, *any* reminder would be helpful, as would any number of possible frameworks for categorizing consciousness. But Chalmers' specific proposal — according to which qualia are quarantined to a class of their own — is (in large part) a direct response to researchers who claim to have made genuine progress in explaining subjective experience, but who have in truth done nothing of the sort. His critique is damning. "It is common to see a paper on consciousness," Chalmers (1995) writes, "begin with an invocation of the mystery of consciousness, noting the strange intangibility and ineffability of subjectivity, and worrying that so far we have no theory of the phenomenon. Here, the topic is clearly the hard problem. [But] in the second half of the paper, the tone becomes more optimistic, and the author's own theory of consciousness is outlined. Upon examination, this theory turns out to be [of one of the easy problems, and] the reader is left

feeling like the victim of a bait and switch" (p. 202). This sort of sleight-of-hand — or simple confusion if it's not deliberately done — creates a muddy swamp of conceptual difficulties whose effect is to retard scientific progress in the study of the mind.

To return to Dennett for a moment, one could argue that his own writings on consciousness suffer from this very weakness. In a review of Dennett's (1991) *Consciousness Explained*, for example, Ned Block (1993) writes that *Consciousness Ignored* would have been a more descriptive title. This is precisely because Dennett manages to float reductive explanations for the "easy" problems only — problems of "access" consciousness, to use Block's (1995) term — whilst simultaneously making a big show of addressing subjectivity (Block's "phenomenal" consciousness — see Appendix). But reductive explanations only work for reducible notions; and Dennett leaves untouched the sort of consciousness whose promised explanation probably sold his book in the first place.

Solving the hard problem

So where do we stand? If what I have argued so far is correct, then the hard problem really is a problem, and it really is hard; and calling attention to it is vital for making progress in the science of consciousness (contra Dennett). But can it be solved? Remember that both Chalmers and I think that an explanation of qualia, if it can possibly be given, will be different in kind from the sort of explanation that is achievable using the tools of cognitive science. Is there another kind of explanation, then, to consider? Chalmers thinks yes: "Given that reductive explanation fails, *nonreductive* explanation," Chalmers writes, "is the natural choice" (1995, p. 209). But what does he mean by nonreductive explanation? I will quote his argument at length here, so that there can be no mistaking his view:

> Although a remarkable number of phenomena have turned out to be explicable wholly in terms of entities simpler than themselves, this is not universal. In physics, it occasionally happens that an entity has to be taken as *fundamental*. Fundamental entities are not explained in terms of anything simpler. Instead one takes them as basic. [For example], in the nineteenth century it turned out that electromagnetic processes could not be explained in terms of the wholly mechanical processes that previous physical theories appealed to, so Maxwell and others introduced electromagnetic charge and electromagnetic forces as new fundamental components of a physical theory. To explain electromagnetism, the ontology of physics had to be expanded. [Accordingly] I suggest that a theory of consciousness should take experience as fundamental (1995, p. 209-10).

Here I think Chalmers is guilty of his own bait and switch. He began by dramatizing the challenge of "explaining" the existence of qualia — and thereby solving the mystery of consciousness — using *reductive mechanics* as the paradigm way to answer "Why does it exist?" for given phenomenon X. He then went on to suggest that reductive mechanics will not suffice when it comes to qualia; and I have argued that he is probably right about that. But when he comes around to giving his own solution to the "hard problem," he elects to hijack that very word "explain" to refer to an act of *positing* — positing, that is, the existence of qualia as fundamental, and hence "expanding the ontology" of not only physics (since qualia are non-physical on Chalmers' account) but indeed the entire world. Why is that a problem? Because to take an entity as fundamental is not to explain it.

Positing, ordinarily understood, is a very different enterprise from explaining. And it is typically much easier to do. To explain something (typically) just *is* to give a reductive or mechanistic account of it. Or, if that cannot be done, at minimum, it is to show why something exists, or is the way it is, by referring to at least one *other* thing and tracing some sort of entailment A to B. To posit, by contrast, requires much less work. You need only to declare, "The thing exists because *it just does*" — a statement which might even be true, but which does not give the sense of ah-hah! that is typically associated with genuine explanation.

Remember, to make progress in the science of the mind, what we need is conceptual clarity. That is why Chalmers' easy vs. hard distinction is so practically useful — contrary to Dennett's view — as I explained in the section above. It is also why baiting and switching and making up new meanings for ordinary terms is such a problem. As consciousness researchers we should always aspire to use terminology in such a way as to maximize (a) humility about what we are

claiming to have shown and consequently (b) genuine understanding in our audiences. Yet by invoking the notion of "nonreductive explanation," Chalmers manages turn what is — on its face — non-reductive, non-*explanatory*, postulation, into something much more impressive-sounding and seemingly more productive than it really is.

So let us take a closer look at Chalmers' proposal: "I suggest that a theory of consciousness should take experience as fundamental" (p. 210). What exactly is the matter with this suggestion? First, we didn't want a theory of "consciousness" in the vague sense which Chalmers, only paragraphs before, had rightly chastised other theorists for using; we wanted a theory of *experience*, in the specific sense of qualia or the what-it-is-like-ness of Nagel (1974). That is what we *mean* by consciousness — it is the "hard problem" — it is the very point of Chalmers' paper. Yet if you reword his proposal using that definition (the interesting definition, the definition worth using) you get the following non-starter: "I suggest that a theory of experience should take experience as fundamental." That doesn't seem quite so fruitful.

Why? Because we didn't want a "theory" of experience either; we wanted an *explanation* of it — maybe not a reductive explanation, if it is impossible to give one, but an explanation nonetheless. We wanted an answer to the question why, in Chalmers' own words, any physical process should "give rise to experience" (1995, p. 208). Taking into account this point, Chalmers' proposal now reads: "an explanation of experience should take experience as fundamental."

This statement is profoundly unhelpful. It saps the word "explanation" of its ordinary meaning; and to borrow from Dennett, it is no "useful contribution to research" at all. This is because, as I want to argue, the hard problem of consciousness involves a hard problem of explanation (hence my title); or as Chalmers himself writes, "*you can't explain conscious experience on the cheap*" (1995, p. 208, italics original). Yet what could be cheaper than "explaining" the existence of qualia by *positing the existence of qualia*? The analogy to physics is actually apt, only not in the way Chalmers intends. Since electromagnetic processes could not be explained in terms of other, simpler, processes, physicists were forced to violate their most cherished maxim of ontological and mathematical parsimony (cf. Clark, 1971), and expand the ontology of physics. Failing to *explain* electromagnetic forces, that is, they *posited* them as basic, and were then able to use them to explain *other* phenomena they observed.

Chalmers is proposing to do the exact same thing with subjective experience. Since he cannot account for it with reference to simpler entities — since he cannot account for it, indeed, with reference to *any* other entity — he takes it as basic instead. That is, he starts with it; he includes it in his premises; he posits it — he does not explain it.

The hard problem of explanation

To explain something, I have suggested, is (at minimum) to given an account of its nature or existence by referring to some *other* existing thing, reductively or otherwise. At least that is what is typically understood by the term when it is used in scientific research. In the case of reductive-mechanic explanations in particular, the answer to 'Why does X exist?' typically involves an appeal to supervenience: the thing exists, whatever it is, because it is either comprised of, or entailed by, some simpler set of extant things. The simpler things *necessitate* the entity of interest, thereby explaining its existence. As the history of scientific discovery makes plain, developing precise accounts of the reductive relationship between higher- and lower-order entities is a painstaking process, and so deserves to be called a "hard problem of explanation" as I have suggested. Indeed, in the case of qualia it may be so hard as to be impossible. That was Chalmers' original point.

What Chalmers attempts instead — in the long passage quoted above — could be termed, then, an "easy" type of explanation: something like "the giving of a theory". This type of explanation (if it deserves the name "explanation" at all) is so easy, in fact, that Chalmers can come up with it from his armchair. To "give a theory" of qualia, as Chalmers does, involves such tasks as situating those qualia in a larger theory, or showing how they relate to other phenomena in a systematic way. Such a theory might certainly tell us (as Chalmers goes on to propose) some inter-

esting things about the *structure* of inner mental life, or the parameters and limits of mind-brain correlation. But at best these are "easy" explanations to the hard problem.

In fact, I would like to go one step further: and here is my appeal for conceptual clarity about explanation. Just as Chalmers has proposed to reserve the word "consciousness" for use in referring to qualia — and to no other aspect of mental life (1995, p. 201) — I would like to propose that we reserve the word "explanation" for use in referring to a special class of statements or arguments. These are statements or arguments which give what Mark Rowlands (2001) calls "epistemic satisfaction" — that sense of "ah-hah!" I alluded to above. As Campbell (2009) puts it, epistemically satisfying explanations include "obvious" cases such as those in which, when presented with an explanation, a person cries out "Eureka!" as a moment of clarity sets in; but also "subtler" cases in which there is "a subjective sense of epistemic progress, a genuine sense of understanding that was absent before an explanation was available" (p. 42).

An illustration may help. Consider this example from Campbell (2009). If you were to say to Oedipus that he should not "attend tomorrow's family reunion because it will be awkward in light of the fact that he slept with Jocasta last night", Oedipus could be forgiven for failing to understand the problem. But if you re-describe "Jocasta" as "your mother," then the situation changes (p. 42). What has changed? Oedipus will have gained a subjective sense of understanding; he will be able to spell out *why* the family reunion would be so terribly uncomfortable**.** In general, explanations will fail when "events or properties are not described in ways that are conducive to the generation of epistemic satisfaction. Hence, although they may be objectively *true* in some sense, they are not *explanatory*, because they do not improve our epistemic situation" (p. 43, emphasis added).

This is exactly what is going on with Chalmers' positing of qualia. His 'theory' might even be correct: it might really be *true* that qualia are fundamental features of the universe, whose existence cannot (even in principle) be explained by referring to any other entity or process. But it is misleading to call this theory an explanation, since we have no more understanding of why qualia should exist after Chalmers' argument as before it — other than by saying "they *just do* exist". That is, we needn't disagree with Chalmers by saying that he is wrong, and that qualia are not in fact basic. They might be. We can simply say that he has failed to *explain* qualia, which is exactly the problem he himself posed in the opening to his paper.

CONCLUSION

Chalmers is right to draw a distinction between the "hard" problem and the "easy problems" of consciousness, and if research is to move forward in this area, scientists and philosophers must be clear about the difference between them. But Chalmers causes problems of his own when claims to have *explained* the relationship between that wet lump of matter in our heads, and the rich qualities of our felt experience. He has done no such thing, on the ordinary understanding of "explain". The hard problem of consciousness — which is a problem of *explanation* — remains unsolved. Yet by reserving the word "explanation" to refer to statements which bring epistemic satisfaction, we may add one lucid drop to the conceptual waters surrounding the scientific study of the mind.

Endnotes

a: A reductive explanation is one that explains some higher-order process in terms of a simpler, more theoretically basic vocabulary taken from a "lower" level of analysis. A full discussion would require a paper of its own; see Neuman (2008) for a useful introduction.

b: A functional explanation is one that explains some phenomenon in terms of what it does. See Levin (2007).

c: I haven't counted, but this is obviously a conservative estimate. As Sloman (2009, 2010) points out, the idea is neither original nor unique to Chalmers, and in a longer paper I would take the time to survey some of its historical formulations. For lack of space, I have included a sampling of highlights in the Appendix.

d: "You need only" makes it sound simple or easy, perhaps. As Chalmers (1995) is careful to point out, however, explanations of this sort — that is, of the reductive, functional, mechanistic variety — are "easy" (with respect to the phenomena in question) only in a relative sense. A full explanation of, say, the focus of attention, might take "a century or two of difficult empirical work" (p. 201) — but at least the explanation is possible. With the hard problem, he contends, such an explanation is not; what is needed is another type of explanation all together.

REFERENCES

Aranyosi I. Chalmers' zombie argument. 2004. Accessed at: http://www.personal.ceu.hu/students/03/Istvan_Aranyosi/Chapter%20IV.pdf

Block N. On a confusion about a function of consciousness. Behav Brain Sci 1995;18:272-278.

Block N. Review of *Consciousness explained* by Daniel C. Dennett. J Phil 1993;90:181-93.

Campbell N. Why we should lower our expectations about the explanatory gap. Theoria 2009;75:34-51.

Chalmers D. Facing up to the problem of consciousness. J Consc Stud 1995;2:3200-3219.

Chalmers D. The conscious mind. Oxford University Press, New York, 1995.

Clark R. W. Einstein: The life and times. World Press, New York, 1971.

Dennett D. Facing backwards on the problem of consciousness. J Consc Stud 1996;3:4-6.

Dennett D. Consciousness explained. Little, Brown, & Co., Boston, 1991.

Howell R, Alter T. Hard problems of consciousness. Scholarpedia 2009;4:4948.

Huxley TH, Youmans WJ. The elements of physiology and hygiene: A textbook for educational institutions. Appleton & Co., New York, 1868.

Jaynes J. Consciousness and the voices of the mind. Can J Psychol 1986;27:128-148.

Koch C. Do not underestimate science. Science 2009;325(5939):392.

Kuijsten M. Close-mindedness and mysticism in science: Commentary on John Smythies's review of reflections on the dawn of consciousness. The Jaynesian 2009;3:1-9.

Levin J. Functionalism. In: The Stanford Encyclopedia of Philosophy. 2009. Accessed at: http://plato.stanford.edu/archives/win2007/entries/functionalism

Nagel T. What is it like to be a bat? Phil Rev 1974;83:435-450.

Newton I. A Letter to Henry Oldenberg. Quoted in: Stanford Encyclopedia of Philosophy: "Panpsychism", 1710.

Neuman Y. Reviving the living: Meaning making in living systems. Studs in Multidisc 2008;6,5-10.

Rowlands M. The nature of consciousness. Routledge, London, 2001.

Ryle G. Thinking and language. Proc Aristotelian Soc 1951;Suppl. 25,68.

Sloman A. Phenomenal and access consciousness and the 'hard' problem: A view from the designer stance. Int J Machine Consc 2010;2:117-169.

Sloman A. Why the 'hard' problem of consciousness is easy and the 'easy' problems hard. Online tutorial presentation. 2009. Accessed at: http://www.worldscinet.com/ijmc/02/0201/free-access/S1793843010000400ref.html

Vimal RLP. Meanings attributed to the term 'consciousness': an overview. J Consc Stud 2009;16:9-27.

Corresponding Author:

Brian D. Earp
Department of Experimental Psychology
University of Oxford
South Parks Road, Oxford, OX1 3UD England
Phone: +44-(0)-206-364-1250
E-mail: brian.earp@gmail.com

Copyright © 2012 by Ass. Crossing Dialogues, Italy

Appendix

Here are some of the more noteworthy, colorful, or interesting formulations of the "hard problem" of consciousness pre-Chalmers (1995):

Isaac Newton (c. 1710): "To determine by what modes or actions light produceth in our minds the phantasm of colour is not so easie"

Gottfried Leibniz (1714): "Perception [is] inexplicable on mechanical grounds, that is to say, by means of figures and motions. And supposing there were a machine, so constructed as to think, feel, and have perception, it might be conceived as increased in size, while keeping the same proportions, so that one might go into it as into a mill. That being so, we should, on examining its interior, find only parts which work one upon another, and never anything by which to explain a perception" (quoted in Aranyosi, 2004, pp. 2-3).

T. H. Huxley (1886): "How it is that anything so remarkable as a [subjective] state of consciousness comes about as a result of irritating nervous tissue, is just as unaccountable as the appearance of the Djinn when Aladdin rubbed his lamp" (p. 170).

Julian Jaynes (1986): "What we have to explain is the contrast, so obvious to a child, between all the inner covert world of imaginings and memories and thoughts and the external public world around us. The theory of evolution beautifully explains the anatomy of the species, but *how* out of mere matter, mere molecules, mutations, anatomies, can you get this rich inner experience that is always accompanying us during the day and in our dreams at night? That is the problem" (p. 128).

I should include here, too, with some explanation, Ned Block's (1995) well-known distinction between "access" consciousness (which involves a type of information processing and is thus what Chalmers would regard as "easy" to explain) and "phenomenal" consciousness (which Chalmers would regard as "hard" to explain). Block's is in fact the earlier paper, and in it he makes the exact point that is most often attributed to Chalmers: "it is of course [phenomenal] consciousness rather than access-consciousness or self-consciousness that has seemed such a scientific mystery (p. 230).