

Reviews

The Character of Consciousness

David J. Chalmers

Oxford University Press, New York, 2010, 624 pp., £65. £18.99

ISBN: 978-0-19-531111-2

ISBN: 978-0-19-531110-5

doi:10.1017/S0031819112000083

Chalmers is a very big name in the philosophy of consciousness and this is a *very* big book about consciousness. Weighing in at over six hundred pages and comprised of fourteen, already published papers (two of which are co-authored), it collects together Chalmers' greatest hits on consciousness in one handy tome. It is comprised of some highly technical and intricate philosophical papers juxtaposed with a couple of more accessible writings that have influenced disciplines outside of philosophy. For anyone wishing to familiarize themselves with the nuances and fine details of Chalmers' approach (or who lacks the full set of his papers already) the book offers excellent value for money. It provides a window on what motivates his approach and, crucially, where his thinking is leading him now.

For aficionados who have been keeping a beady eye on these developments – what's new? Apart from a sprinkling of authorial comments within the various chapters, comments that serve to highlight links between them, there is also the book's seventeen-page introduction. Its overview gives Chalmers' own assessment of how the structure of the chapters interrelate with one another; flags up the important new ideas and proposals; and highlights the respects in which his thinking has moved on since the publication of *The Conscious Mind*.

While the book's title might suggest that consciousness has a singular nature, Chalmers is clear that, for him, it has multifaceted characters, many interesting dimensions. Thus any attempt to understand it requires concerted effort along multiple and inter-connected fronts. Unavoidably, any fully convincing account of it must address questions about: (1) the problem that phenomenal consciousness poses for standard materialism (chapter 1); (2) how it can be studied scientifically in the light of that problem (chapters 2 to 4); (3) its metaphysical nature; (4) its epistemic character and how that relates to the nature and acquisition of phenomenal concepts (chapters 8 to 10); (5) its intentional and representational character, and the implications this has for how we understand the nature of the external world (chapters 11 to 13); and (6) how it can exhibit both unity and diversity.

Any full inquiry into the nature of consciousness ought to have something to say about each and all of these several dimensions – at least. Doing justice to these topics demands separate and detailed investigations. Still in any satisfactory overall account the conclusions of such investigations should all hang together coherently. The aim of providing just such a workable overall account is the rationale behind Chalmers' grand systematizing approach to this subject matter. Even if one rejects the final product, one cannot fail to be impressed by the care, attention and cunning argumentation deployed in its construction. Chalmers is thorough. He attempts to give an exhaustive and even handed consideration to every possibility – or at least every possibility that he takes to be a serious contender – in order to persuade the reader by cool headed argument.

This approach explains the book's enormous size. It is also reflected in the style of the writing. While some chapters – for example, chapter 1 – are accessible to non-analytic philosophers, most chapters offer extraordinarily refined analyses of the topics under investigation. These analyses will no doubt serve as key reference points for professionals working in analytic philosophy of mind. Yet even though the positions are set out carefully, even professional philosophers will need to invest some time and care if they are to get a full and accurate grasp of Chalmers' overall account and to keep track of his particular reasons for advocating it. Most chapters make free use of technical jargon – e.g. distinguishing between primary and secondary intensions, positive and negative conceivability – invoked in order to make the conceptual possibilities and argumentative structure more perspicuous. While in some cases this may be necessary and even helpful, it does not make for an easy read. For anyone unfamiliar with this style of philosophy, be warned: reading this book is not a task for the light headed or faint hearted.

The driving insight behind Chalmers' approach hasn't changed; his fundamental challenge to the field of consciousness studies is that all standard forms of materialism are false. From the fact that it is impossible to understand or reductively explain the phenomenal in terms of the physical broadly conceived, it follows from this that there is more to reality than is dreamt of, or posited by, modern day materialists or physicalists. Those who hope for explanatory or ontological reductions using only the resources of the physical sciences and what they pick out are out of luck. To accept that phenomenal consciousness exists requires recognition of 'extra ingredients' in reality above and beyond physical properties. Those hoping to get by with less are confronted with the hard problem of consciousness, a problem inspired by the alleged existence of an explanatory

Reviews

lacuna – the infamous explanatory gap that cannot be closed by any account in terms of structures, functions or dynamics, no matter how subtle, fine-grained and wide ranging the latter are. Such phenomena are, at best, nomologically associated with experiences of phenomenal consciousness. Thus it may be that in our world they are related to it and required for it as a matter of natural law but they are not logically related to it or required for it.

A great deal of this new book is devoted to explicating the nature of the hard problem and attempting to show that there is no way to deny or overcome it. Its opening chapter (which hails from 1995) sets the stage. It argues that when it comes to consciousness, explanations of the sort that can be provided by the natural sciences always and necessarily leave untouched what most needs explaining – i.e. the properties of what it is like to have a given experience. Phenomenal properties are intrinsic properties of experiences – for this reason they are systematically overlooked by, and cannot be accounted for in terms of, or reduced to functional, structural or dynamic properties. Any attempt to explain the phenomenal aspects of experience by appeal to properties of that sort always fall short of what is required – such explanations can only provide more detail about structure and dynamics – they never get at, never even mention, the requisite phenomenal features as such. Intuitively, the problem is illustrated by the fact that ‘for any complex macroscopic structural or dynamical description of a system, one can conceive of that description’s being satisfied without consciousness’ (122).

Chalmers makes an interesting admission concerning this line of argument in the book’s introduction. He claims that this simple argument, concerning explanation, is ‘in a sense’ more fundamental than conceivability arguments that have attracted the most philosophical attention (xv). Without denying their important uses, nor shying away from assessing them in great technical detail, at length (in part III), he denies that anti-reductive arguments depend essentially on appeal to standard philosophical thought experiments about zombies or what Mary learns on leaving her confinement. This, however, understates the philosophical complexities of the situation. Appeal to explanatory considerations is no doubt more intuitively compelling and more persuasive to those outside of or unused to philosophy. But acknowledging that an explanatory gap exists depends on one’s background philosophical assumptions – assumptions that are exposed and tested by reviewing one’s reactions to the thought experiments. So it is not clear in what sense the explanatory argument could be more fundamental. Indeed, although it is to be hoped for, it is not clear that there is any Archimedean point for fashioning

conclusive arguments in this domain – one that all parties can agree upon, that doesn't beg any crucial questions. For more on Chalmers' response to this, see his revealing discussion on pages 33–34.

Some, for example, do not accept Chalmers' framing assumption. This is the most obdurate response to the hard problem – that of Type A materialists. They deny the existence of any epistemic, metaphysical or explanatory gap. Phenomenal zombies are inconceivable and, even locked in her black and white room, there is nothing that Mary doesn't know about what it is like to experience redness. Although this needn't seem obvious, a wholly third person approach – one that may involve complete knowledge of a correct physics – provides all that is required for understanding the phenomenal. Put this way, Type C materialism collapses into Type A materialism (as Chalmers himself argues). Type Cers hold that phenomenal zombies only seem conceivable to us because of our current state knowledge, they are not conceivable 'in the limit'. Type C types are really Type A types who add a 'the cheque is in the post' rider to their account.

By far the most popular response to Chalmers' challenge has been to adopt some sort of Type B materialism – to accept that there is an epistemic and explanatory gap while denying the existence of an ontological gap. Phenomenal zombies are conceivable, but this is put down to peculiar features of phenomenal concepts; nothing interesting follows from the fact that we conceive such monsters about the metaphysics of consciousness. This is to opt for a non-explanatory physicalism. As a result, for Type Bers, the link between the physical and the phenomenal is thought to be 'epistemically primitive'; one cannot deduce phenomenal truths *a priori* from known physical truths. That doesn't preclude physical facts being all the facts. The phenomenal might just be the physical differently described – under a different, irreducible, guise or mode of presentation. One might hold this while postulating identities, as Papineau does, and arguing that there is no need to explain them. Does it make sense, after all, to ask *why* Clark Kent is Superman?

No doubt due to its popularity, the book focuses mostly on undermining this sort of approach, providing a range of extensive and detailed arguments against its viability. For one thing, it is complained that an appeal to unexplained identities won't do since identities in other domains can be deduced from more basic truths (see 117). Other arguments, based on Chalmers' own, controversial, two-dimensional semantic framework, are also brought to bear to show that there is no way to avoid the link from ideal conceivability to metaphysical possibility (see chapter 6 and appendix). Even more

Reviews

powerfully, building on his analysis of the nature and formation of phenomenal concepts (chapter 8), and the implications this has for epistemology (chapter 9), Chalmers advances his master argument against Type B materialists who invoke the phenomenal concept strategy (chapter 10). The dilemma he poses is, crudely but essentially, that Type Bers must either assume that it is (a) possible to explain why there are phenomenal concepts that generate the explanatory gap, given our makeup, in physical, non-phenomenal terms or (b) it is not. If they assume (a), their position collapses into Type A materialism. If they assume (b), they lack a means of explaining the epistemic gap in tune with physicalism. Either way the phenomenal concept strategy affords no escape from the hard problem.

Chalmers concludes that the only live options for responding to the hard problem are adoption of some form of Type D interactive dualism, Type E epiphenomenal dualism, or Type F Russellian neutral monism. In the introduction he confesses to having less sympathy for epiphenomenalism than he once did. Having myself promoted the fortunes of a non-Russellian monism, one that has closer affinities to Bradley's Absolute Idealism, I find Chalmers' final score sheet encouraging. A non-physicalist monist can allow that phenomenal experiences admit of true physical descriptions, without assuming that such descriptions exhaust or fully characterize all aspects of their nature. Experiences can be enjoyed and truly described in ways that do not invoke the vocabulary of physics. When events are enjoyed or described in those different ways, different features are highlighted. There is no ultimately privileged way for describing experiences that captures all of their essential properties. However perfect physical science may become its descriptions and explanations will always be as conceptually limited and interest-sensitive as any other. Even ideal physical descriptions do not provide a transparent window into every essential aspect of what they pick out in extension. Chalmers flirts with a similar sentiment, accepting that 'perception and even science do not reveal the entire intrinsic character of the world' (xxvi). However, in chapter 13, which he regards as his favourite, he develops this thought in an un-idealist way, by advancing the matrix hypothesis in the form of a metaphysical hypothesis.

Taking any one of these routes, by themselves, does nothing positive to clear the air of the explanatory mystery associated with the hard problem – where this is understood as requiring an intelligible explanation of how the relation between the phenomenal and the physical. Still Chalmers hopes to provide a positive proposal to

address it by advancing a different sort of science of consciousness. He envisions a science of consciousness that is essentially correlative rather than reductive. Yet, while he recognizes that robust scientific research into consciousness has blossomed during the past sixteen years, he confesses to only limited progress on finding bridging principles to link first and third personal phenomena – as he construes them. Nevertheless, ‘if one squints one can discern the possibility of framework on the far horizon’ (xvii).

It is arguable that a problem with Chalmers’ way of understanding the nature of a future science of consciousness is that in seeking bridging laws it is too closely associated with assumptions of naturalistic dualism as opposed to some variant of monism. On some variants of the Type F monistic reply there simply is no relation here that needs explaining. From this perspective solving the hard problem isn’t just hard – it is impossible. The articulation of the hard problem serves the important function of exposing flaws in materialist thinking. It shows that we need to rethink the nature of phenomenal consciousness, in a quite fundamental way. Nevertheless, it is to exercise practical wisdom to avoid trying to give a straight solution to it. This is not to cop out. It’s a waste of time to try to solve problems that cannot be solved. The really hard problem is getting certain philosophers to see the hard problem is an impossible problem.

What of other aspects of Chalmers’ positive project? He holds that intentionality and consciousness are ‘intimately connected’; they are ‘intertwined all the way down to the ground’ (371). As such any satisfactory accounts of these – the most important phenomena of mind – must speak to another. In a bid to provide such an account, Chalmers sets out the stall for an unusual sort of representationalism about phenomenal properties – defending a nonreductive, impure, narrow, Fregean version of representationalism (341). Unsurprisingly, given his arguments against reductionism, he bucks the trend of seeking to understand phenomenal properties in terms of representational properties, allying himself more closely with those that hold there are at least some important kinds of representational properties that might be grounded in phenomenal properties.

He defends the view that experiences have phenomenal content – that they represent the world as being a certain way – a way that is assessable for accuracy; thus they have a specific kind of representational content with built-in conditions of satisfaction. This, he admits, is a ‘substantive thesis’ (383), but it is one – he thinks – there is good reason to accept. In his view experiences enjoy a veritable smorgasbord of content – there is Russellian, Fregean and

Reviews

Edenic content. As he presents things, such pluralism looks unproblematic because he characterizes Russellian contents as extensions and Fregean contents – modes of presentation – are conditions on extensions. Ultimately, due to problems with Russellian representationalism he defends the view that phenomenal representational content is a kind of Fregean content – a mode of presentation with a content captured by something like the following condition of satisfaction: the object or property that normally causes experiences of such and such a phenomenal type. Yet he recognizes that Fregean modes of presentation, so defined, lack phenomenological adequacy and threaten to overintellectualize the phenomenal contents of experience (a concern that lingers even if he can defend the idea that such contents are nonconceptual, as is his wont – see pages 368–9).

In response, he posits yet another kind of content – Edenic content – to help refine and supplement his account of Fregean content. Edenic content is how things appear to us phenomenologically – how they would have to be if our experiences were to be perfectly veridical. But on the assumption that we don't live in Eden, there is a gap between how things seem and how they are. To use Chalmers' example – we can, and often do, experience perfect redness even though our world at best instantiates imperfect redness – where the latter might be some physical property, such as surface spectral reflectance.

Does this mean our phenomenal Edenic content always and systematically misrepresents? Chalmers denies this. For he holds that imperfect redness will match perfect redness if it plays the role that perfect redness would have played in Eden – i.e. that of bringing about phenomenally red experiences. With the less stringent notion of imperfect redness and the notion of matching in play, Chalmers aims to secure the idea that our representations are often veridical in an ordinary sense. We can accept this if we are not overly demanding about what veridicality requires – that is, if we require less than what would be required for veridicality in Eden. With this apparatus in play, the Fregean conditions of satisfaction can be formulated in a revised way to capture what is represented by phenomenal contents. For example, 'a phenomenally red experience will be imperfectly veridical iff its object has the property that normally causes phenomenally red experiences' (403). This representational content is plausibly often satisfied, even in our imperfect world. Thus 'Our experience presents an Edenic world and thereby represents an ordinary world' (406). Since the ordinary Fregean and Russellian content of experiences are derived from Edenic content, it turns out that 'Edenic content is that key' (454).

This account is ultimately not very satisfying. For Chalmers, ‘It seems intuitively clear that perceptual phenomenology, by its very nature, involves representation of the external world’ (344). In line with this intuition, he holds it is ‘not implausible that there is something about consciousness that by its very nature yields representation of the world’ (371). Yet claiming that phenomenality and representation are linked in very basic and internal ways does nothing to help us understand how experiential states of mind could possibly make inherent contentful demands on the world. It is not at all clear how simply enjoying a certain qualitative experience – say, the redness of red – could suffice to represent the world as being a certain way, where this entails its having a kind of content that makes a claim on reality. It is not clear how qualitative experiences – by themselves – could ‘attribute’ certain properties to reality or state which conditions the world would be required to satisfy if one’s experience were to be accurate or veridical – not even in the ordinary sense.

A more robust account than Chalmers offers is required if we are to make intelligible how experiences might intrinsically possess such properly contentful properties. Chalmers offers no such explanation, nor does he pretend to. While this is wholly in line with his promotion of non-reductionism, in the end the lack of explanatory power makes his account more questionable, not more attractive. This is revealed, for example, when he asks, rhetorically, how mental states come by Edenic content, given its pivotal importance in his account. Disappointingly but tellingly, he declines to answer, noting only that, ‘We do not yet have a good theory of how mental states represent properties at all’ (418).

That is surely true. All existing attempts to explain content as naturally occurring, purely physical, information processing or bio-functional phenomena have proved inadequate. Indeed, there are excellent reasons to suspect that mental content is logically distinct from, and irreducible to, properties, relations and functions identified by the natural sciences. A plausible diagnosis of the source of these failures suggests that there is a hard problem of content – one that is no less serious than the hard problem of consciousness. But, if so, perhaps rather than sticking with certain fashionable intuitions about the representational nature of perceptual experience, and building elaborate and intricate, but ultimately non-explanatory frameworks to accommodate such intuitions, we might also revisit standard thinking about these topics too. For all that Chalmers offers us his magnificent, mammoth and challenging *tour de force*, there are many conservative strands in his thought that warrant serious

Reviews

challenge. The first step is to get to know his arguments in detail – and there is no better place to start than with this book.

Daniel D. Hutto
d.d.hutto@herts.ac.uk

Love's Vision

By Troy Jollimore

Princeton University Press (Princeton, Oxford), 2011, 197 pp., £24.95

ISBN: 9780691148724

doi:10.1017/S0031819112000095

Plato spoke of love in terms of vision, and Jollimore grants also that love can involve a truthful vision of reality, though he rejects Plato's inference that our vision is blinkered if love remains fixated on human beings. Jollimore is concerned exclusively with love for and between persons, and the love in question comprises erotic love and friendship. These loves, he says, 'can be treated by a single account' (xiv), and this 'friend/lover model' is most relevant to love *between* human beings. *Agape* is excluded because it is doubtful whether we are capable of feeling it (xiii).

Love is a vision in the sense that it involves 'an appreciative attention directed towards the beloved's positive qualities' (7). The claim contravenes the common refrain that love is blind, a refrain which tends to go hand in hand with the idea that the lover invents rather than discovers the qualities she finds in her beloved. Jollimore allows that some beloveds bear no resemblance to their lover's image of them, but resists the assumption that such pathological cases are paradigmatic. So genuine love is not blind in *this* sense, but it might be thought nonetheless to involve misperception and delusion about the *value* of the beloved's qualities – everything about her becomes wonderful, even her toothmarks on a pencil (8). Hence the suspicion that love is fundamentally irrational.

Jollimore resists this conclusion and argues that love can be justified by reasons, the most powerful sources of reasons being 'the attractive and otherwise valuable properties of the beloved' (13). He is quick to point out, however, that awareness of such properties does not rationally obligate the lover, or anyone else for that matter, to fall in love with their possessor: 'no list of attractive or otherwise valuable features of B is such that a person cannot admit that B possesses all of those features and yet fail to love B without being irrational' (14). Likewise, there is no requirement that the lover fall