

Two Dogmas of Reductionism: On the Irreducibility of Self-Consciousness and the Impossibility of Neurophilosophy

*By Joseph Thompson**

Two fundamental assumptions have become dogma in contemporary Anglo-American philosophy of consciousness: that everything about consciousness can be explained in physical terms, and that neuroscience provides the uniquely authoritative methodology for approaching the essential questions. But there has never yet been a successful physical explanation of subjective first-person experience, and reductionism fails to account adequately for thought, reason, and a full range of objects proper to philosophy. Tracking the deep divide within the analytic tradition, I bring a ‘continental’ (German) perspective to bear on recent work from Nagel and Chalmers which shows the reductionist neuroscientific agenda to be incapable of completion, for systematic reasons. Physicalism can explain only structure, function, and mechanism; but self-consciousness is always already embodied and embedded in multiple contexts beyond the structures and functions of brain activity. Consciousness needs to account for itself, to itself, on the terms in which it experiences itself. No explanation of the form provided by ‘neurophilosophy’ is adequate to the most fundamental and essential phenomena of self-consciousness, and neurophilosophy can never philosophically explain or justify itself on its own terms and by its own methodology. These are insuperable limitations for any explanation aspiring to be comprehensive, and such problems have brought contemporary antireductionists in the English-speaking world back around to positions which strongly resemble the ontology and phenomenology of German-language philosophers, particularly Kant, Hegel, and Heidegger.

Two Dogmas of Reductionism

There is a deep division within contemporary analytic (Anglo-American) philosophy of mind and consciousness, between *reductionists* and *antireductionists*: those who think there are only physical problems to be solved, and those who maintain that any physicalist explanation will be fundamentally inadequate to a full account of consciousness and self-consciousness. From the standpoint of German philosophy, I bring a

*Associate Professor, University of Alaska Fairbanks, USA.

‘continental’ philosophical perspective to bear on this dispute, aiming to clarify and sharpen recent developments, strongly supporting the antireductionist position, and outlining a critique of what is called ‘neurophilosophy.’

In the philosophy of mind and across the neurosciences, consciousness reductionism in its various forms has been the dominant approach for the past three decades. Consciousness, like everything else, like all there is, is ultimately to be explained in physical terms: ‘reductive materialism is widely assumed to be the only serious possibility.’ The physicalist version of reductionism presents, not arguments, but only ‘hope’: ‘the hope that everything can be accounted for at the most basic level by the physical sciences, extended to include biology’ (Nagel, 2012). In the neuroscientific version of consciousness reductionism, the privileged level of description is neurobiology and the cognitive sciences. Consciousness is to be accounted for as neural activity, which is supposed to afford unique access in answering questions in the philosophy of mind.

Whether in its physicalist or neuroscientific forms, *reductionism* has a negative connotation, suggesting the dogmatic generalization that everything there is can be reduced to purely material, physical terms—to molecules in motion and nothing more. With respect to consciousness and self-consciousness specifically, reductionism purports to explain these phenomena entirely in terms of the physical processes, structures and functions of the brain: the neuroscientific replacement of the concept of ‘mind’ or ‘self’ by neural activity and nothing more.

From the confident tone of such prominent figures as Dennett and Churchland, one would get the impression that the reductionist agenda has largely succeeded, and that the big questions of consciousness have been essentially answered. Even if all the explanatory details have not yet been worked out, we can be sure that they will be, and we can know in advance what form those answers and explanations will take. The outdated problems, approaches, and language of what used to be called ‘philosophy of mind’ are being replaced by the methodology and corrected terminology of neurophilosophy, the philosophical application of the empirical results of neuroscience.

On the other side of the divide, however, for more than twenty years there have been forceful arguments against reductionism, culminating in the most recent works of Nagel and Chalmers, both from Oxford University Press: a closely argued, nearly 600-page exhaustive account from Chalmers in 2010, *The Character of Consciousness*, and a masterfully brief tour de force on *Mind and Cosmos* from Nagel in 2012, with the provocative subtitle *Why the materialist neo-Darwinian conception of nature is almost certainly false*.

The critique leveled by Chalmers and Nagel shows the attempt to explain consciousness in purely material, physical, or neuroscientific terms to be not only incomplete but incapable in principle of completion. There are strong, systematic reasons relating to the inherent limitations of physical explanation, why any such account must fail to address what Chalmers (1997) has made famous as ‘the hard problem of consciousness.’ For any assumption to the

contrary, the burden of proof is now on the other side; but no compelling counterarguments or rebuttals have been produced, and it seems doubtful whether any could be.

We emphatically do *not* know that consciousness *must* be physical. That is neither a law of logic nor an empirical fact. And indeed both sides recognize that the subjective experience of self-consciousness has never yet actually been captured by any physical, materialist reduction. No science has been capable of ‘explaining how and why consciousness arises from physical processes in the brain’ (Chalmers, 1997). Indeed the phenomena most in need of explanation cannot be addressed by any reductionist account, due both to the nature of subjective experience, and to the nature of physicalist, reductionist explanation itself. Despite general acceptance within the sciences and among many scientifically minded philosophers, reductionism is neither self-evident nor unproblematic. For Nagel it is highly doubtful: ‘I realize that such doubts will strike many people as outrageous, but that is because almost everyone in our secular culture has been browbeaten into regarding the reductive research program as sacrosanct, on the ground that anything else would not be science’ (2012).

Neurophilosophy and consciousness reductionism are thus predicated upon two highly problematic assumptions: that everything (including the phenomena of consciousness) can be explained in physical terms; and that neuroscience provides the only reliable and accurate methodology for knowing things about consciousness. Neither of these suppositions has ever been determined, whether by scientific evidence or by philosophical argumentation. Numerous objections stand against these assumptions. Lacking the requisite philosophical justification, they may be identified (following Quine) as ‘two dogmas of reductionism.’

Consciousness Not Explained

Faced with the task of explaining consciousness, as Dennett purports to do, one of the most immediate and important issues obviously concerns the *explananda*: what exactly needs to be explained and accounted for? This is a fundamental question that directly informs the way the entire explanatory process is undertaken, and how the question is answered will dictate what sort of results are counted as successful explanations. In a work with the ambitious title *Consciousness Explained*, Dennett (1991) assures us he will ‘explain the various phenomena that compose what we call consciousness, showing how they are all physical effects of the brain’s activities’—but he never does make good on that claim. As he himself must admit:

Adopting materialism does not by itself dissolve the puzzles about consciousness, nor do they fall to any straightforward inferences from brain science. Somehow the brain must be the mind, but unless we can come to see in some detail how this is possible, our materialism will

not explain consciousness, but only promise to, some sweet day
(1991).

More than two decades later reductionists are still at the point of insisting upon that ‘somehow,’ and still a long way off from that ‘some sweet day.’ For a philosophical theory of consciousness it simply begs the question, to assert that the only things to be explained are those functions and structures of the brain amenable to neuroscience and physicalist reduction. Reductionism as a philosophical presupposition remains a dogma: it starts from the controversial assumption of materialist metaphysics, with everything that entails. Meanwhile there is overwhelming phenomenological evidence to the contrary that must be explained away in order for reductionism to work. Any adequate account of consciousness must no doubt be informed and supplemented by a thorough understanding of the physical mechanisms of the brain, but a physicalist reduction would definitely be no substitute for a comprehensive philosophical theory, not least because it would fail to address or justify its own fundamental principles and presuppositions.

Ontological and epistemological assumptions about what kinds of things exist and need to be explained cannot themselves be grounded by any appeal to neuroscience: these claims fall within the domain of philosophy proper, and as such they require philosophical argumentation. Chalmers (1997) notes that ‘to establish *this* position—that there really is nothing else to explain—one might think that extraordinarily strong arguments are needed.’ Yet the proponents of this view have provided no such arguments: rather it has been taken as uncontroversial, even axiomatic. But we cannot know *a priori* that the philosophical explanation of consciousness will take the form of a physicalist reduction, whether at the neurochemical or quantum level or otherwise, lacking (as we do) the ‘extraordinarily strong arguments’ required for such a high-level determination.

Not only has no such explanation been provided, but there *can be* none in these physicalist and neuroscientific terms—none that actually accounts for the phenomena in question, rather than sidestepping the real problem and dealing with a different one instead, or denying that there is a problem at all. Reductionism assumes that the only possible objects of knowledge about consciousness are physical and brain-functional, and from this epistemological presupposition draws the further and much stronger *ontological* conclusion, that what is not explicable in terms of neuroscience and biology, chemistry and physics, somehow *does not exist* or needs no explanation. That is not a scientific truth, nor could it be established by any science, not even in principle. The claim that scientifically derived empirical facts are all that we can know about consciousness is not a scientifically derived empirical fact. As Chalmers (1997) responds to his critics:

Proponents of the ‘no problem’ view sometimes like to suggest that their view is supported by the results of modern science, but all the science that I know is quite neutral here: I have never seen any

experimental result that implies that functions are all that need to be explained. Rather, this view seems to be rooted in a philosophical claim...

The *subject* of first-person mental life-as experienced pre-philosophically, given prior to any analysis-must be central among the default explananda in the philosophy of consciousness, until and unless it is proven (as it never yet has been) that the subjective selfhood we experience is in fact an illusion. Otherwise that remains an extremely counterintuitive (perhaps even incoherent) position, directly contradicted by immediate evidence, to which one is driven only because the assumptions of physicalism require such a denial. For the antireductionists, far more plausibly and practicably, ‘conscious subjects and their mental lives are inescapable components of reality not describable by the physical sciences.’ These realities cannot be avoided or dismissed by any account of consciousness that aspires to be anything close to comprehensive: conscious minds ‘are among the data that a theory of the world and our place in it has to explain’ (Nagel, 2012).

Accounting for the Nonphysical

It has long been argued in at least one main line of Western philosophy—and in numerous nonwestern traditions—that not everything which has reality and significance is physical. From Pythagoras and Plato to Descartes and Kant, the rationalist tradition has always had powerful arguments for the existence and reality of the nonphysical. Still, the intractable problems of mind-body interaction have led empirically oriented philosophers to try to avoid dualism at all costs (Dennett, 1991). But it does come at a cost: anything apparently not physical must be *epiphenomenal*, essentially denied reality. This encompasses all evidently nonphysical features of self-consciousness, subjective experience, thought, rationality, language, ‘the mind’ itself. It is no coincidence that Plato and Descartes are among the least reputable in contemporary neurophilosophy. But reductionists remain confronted with obvious problems of ontological plurality, even if dualistic minds do not exist but only brains. Reductionism is already incapable of accounting for the indisputable actuality of mathematics and logic: both are clearly related to the physical, but also retain their formal relations and truth apart from and independent of particular concrete physical expressions or examples. Math and logic have objects which appear to be distinctly nonphysical, nonmaterial, and to have a pronounced universal character: they go beyond the physical, and thereby call into question the explanatory range and power of any physicalist reduction. There is an unreconciled dualism, *a priori*, at the heart of any supposed materialist monism.

Reason itself is another central philosophical concern, related to math and logic, which does not appear to be physically reducible, but which is obviously operative and capable of generating truth and knowledge of all kinds, some

‘pure’ (in Kant’s sense) and some ‘practical.’ Reason must be accounted for on its own terms, for only with reason can we explain, correct, or answer anything. Reason is its own ‘final court of appeal’; it has ‘completely general validity, rather than merely local utility’ (Nagel, 2012). For ‘practical’ purposes we must be realists about reason: the rational is actual, and rational relations obtain independently of any brain activity. Reason is one of a whole broad set of realities that we experience and need to explain: Nagel (2012) includes ‘consciousness, intentionality, meaning, purpose, thought, and value’ along with reason, as chief among the explananda incapable of physicalist reduction, for which a higher-order level of ontological description is required.

If we add to that list other essential philosophical concerns like language, art, religion, literature, history, culture, ethics, politics, justice—it becomes more and more apparent that no scientific reduction is going to be able to address and explain all these, and most certainly not a neuroscientific reduction. But all these areas of inquiry are the rightful domain of philosophy: they are part of the broad set of interconnected questions that define philosophy, and they show just how limited neurophilosophy must be considered in comparison with philosophy proper.

The empiricists have never been able to account for conceptual reality empirically, not in any convincing way, because pure forms of reasoning and inference are inherently abstract and *a priori*. The intelligible realities of mathematics and logic cannot be denied actuality, yet whatever ‘being’ they have must be other than physical. They may have physical *instantiations*, but *what* is thereby instantiated is *conceptual, universal*, and any adequate ontology must be able to explain the fundamental reality of reason, logic, abstract concepts, universals. The very existence of logical necessity and the fact that true propositions can be deduced from other propositions is itself something that must be accounted for. Any and all empirical arguments must presuppose *a priori* principles of reason, whose justification is not biological, and which are not reducible to the terms of physics or any natural science. If reason, math, and logic are nonphysical, then not everything is physical—so then not everything can be explained in physical terms.

Math and logic present confounding ontological counterexamples, if not a flat-out refutation of physicalist, materialist monism. They reveal the insuperable limitations of reductionist explanation: dualism is already an inescapable problem, before the even harder problem of subjective experience and self-consciousness is factored in. If physical reductionism cannot account ontologically for the abstract, nonphysical and universal truths and objects of math and logic, *a fortiori* it cannot explain the forms of rational consciousness which apprehend those truths and objects. So now there are at least *three* fundamental ontological categories to account for: (1) the conceptual (e.g. mathematics), no less than (2) the material/physical, and (3) self-consciousness. It is striking how directly these categories recall the Hegelian triad of Logic, Nature, and Mind [*Geist*]: we appear to have circled back around at a deep level to German philosophy, as neo-Kantians, neo-Hegelians and neo-Heideggerians.

Consciousness, thought, reason, math and logic join what is a long list of ontologically significant phenomena whose being cannot be explained or accounted for in purely physicalist or materialist terms. Not only mathematical and logical entities are nonphysical; so too are any abstract concepts, including those of ethics and value: persons, rights, equality, dignity, fairness, justice, ‘the good, the true, and the beautiful’—these all go beyond the merely physical substrate. This is not to say that they must be ‘supernatural,’ in the derogatory or reproachful sense of the word. It means that they cannot be explained in physical or neuroscientific terms: an expanded set of ontological categories and an expanded conception of ‘nature’ are necessary to overcome the explanatory gap. To adequately account for the phenomena of self-consciousness, *phenomenology* is necessary: as Chalmers (1997) argues, phenomenological approaches to the ‘hard problem’ must be ‘absolutely central to an adequate science of consciousness: after all, it is our own phenomenology that provides the data to be explained!’ Any Anglo-American attempt at explaining consciousness must reckon with the critique already advanced by the German philosophical tradition, concerning the foundations and general features of the phenomenological approach: not only in Heidegger’s or in Husserl’s sense, but primarily along the lines of *Hegel*, that original post-Kantian phenomenologist of consciousness.

The Impossibility of Neurophilosophy

The neurophilosophical project was therefore bound to fall short, because neuroscience cannot substitute for phenomenology or ontology: it cannot address or explain entire areas of philosophical inquiry. No science has ever scientifically established its own philosophical presuppositions: the sciences do not empirically justify the logical, ontological, and epistemological determinations which are the conditions for the possibility of science.

Of course the philosophy of nature has branched off over the centuries into many sciences; but science has never been able to ‘replace’ philosophy, and there is no reason to expect that neuroscience will fare any better. Philosophy and science have a great deal in common, and there is a long and likely familiar set of arguments for their similarity and affinity that need not be rehearsed here. But there are also numerous points differentiating philosophy from science, which show how and why they are not the same thing; science on its own cannot establish philosophical conclusions any more than philosophy on its own can establish scientific propositions. But philosophy, crucially—as Kant and Hegel practice it, for example—can *account for itself* on its own terms, while the sciences cannot.

From the overarching view pejoratively termed ‘scientism’—so called because it dogmatically insists that the sciences can explain everything—we can coin the term *neuroscientism*, to designate the dogmatic extension of neuroscientific reductionism into other domains of philosophy, the attempt to appropriate and reframe the ‘traditional’ problems of philosophy into questions

that can be addressed by neuroscience. *Neurophilosophy* is meant to be a *science*, as Churchland (1989) makes clear in her subtitle: *'Toward a unified science of the mind/brain.'* More recently, Churchland (2008) claims that traditional philosophy of mind has been taken over and 'replaced' by neuroscience:

Since the weight of evidence indicates that mental processes actually are processes of the brain, Descartes' problem has disappeared. The classical mind/body problem has been replaced with a range of questions: what brain mechanisms explain learning, decision making, self-deception, and so on. The replacement for 'the mind-body problem' is not a single problem; it is the vast research program of cognitive neuroscience.

But 'the really hard problem of consciousness,' according to Chalmers (1997), 'is the problem of *experience*,' and it has assuredly not 'disappeared': 'any neurobiological or cognitive account will be *incomplete*, so something more is needed for a solution to the hard problem.' For Nagel (2012) 'the mind-body problem is not just a local problem': far from having disappeared, 'it invades our understanding of the entire cosmos and its history.' The implications are profound, and the explanatory task that much greater, when the 'really hard problem' takes us from neurons out to 'the entire cosmos and its history.'

Few neuroscientists would presume to make pronouncements concerning the cosmos. Neurophilosophy is not in a position to give a comprehensive analysis of consciousness in its cosmological context, considering its full implications. Even the object directly at hand remains a 'really hard problem,' an *unsolved* problem: showing how and why the mind is nothing more than brain activity. But the 'impossibility of neurophilosophy' refers more broadly to the inability of cognitive neuroscience to give an adequate account of fundamental philosophical issues. It does not and could not address entire branches of foundational questions on its own terms. Neuroscience cannot provide a systematic and coherent integration. As a philosophical methodology it is inapplicable and incapable outside its area; it certainly cannot presume to provide an account of the totality. In these respects neurophilosophy is and must remain limited and incomplete, due to inherent differences between science and philosophy proper.

Scientists have many times tried to claim epistemic priority for science over philosophy, but such claims are themselves philosophical and therefore self-undermining. Science has always been closely related to philosophy, but it has never been adequate to 'replace' it, for systematic reasons relating to the kinds of empirical questions settled by the sciences: causal questions, physical, material, functional questions. Their answers are framed as scientific explanations or hypotheses. But those are not the forms of fundamental philosophical questions and answers, a difference which is essential to both philosophy and science. The normative claims of philosophy (including the

priority given to scientific knowledge) are not like the empirical propositions of the sciences themselves, and they are established in very different ways. The broad categories of philosophy and science are as obviously distinct as they are also in subtler ways similar. A scientific fact discovered e.g. by functional magnetic resonance imaging (fMRI) does not answer a philosophical question; and conversely if it is a philosophical question, then the answer is not discoverable by laboratory experiments or fMRI scans. Among other things, philosophy asks *metaquestions*: metaethical or metalogical or metaphysical questions are not settled by fMRI or by describing how brain mechanisms work.

Neurophilosophy therefore cannot be considered philosophy proper, but only a branching off from some areas of inquiry in philosophy, while leaving other entire areas completely open and untouched. Neurophilosophy provides essentially no account, for example, of logic or ontology, and no ethics or ethical theory that is not already question-begging. Its restricted range offers explanations for only a subset of issues within the philosophy of consciousness. The entire orientation and approach appears excessively ‘neurocentric’: overly privileging the brain, yielding an ontologically incomplete set of phenomena as compared to the full range of objects of philosophical investigation. It is evident that fundamental questions, metaquestions which are properly philosophical, could never be settled by an appeal to neuroscience or indeed any science. There is no scientific method to ground, validate, or verify metaethical determinations: rather such evaluative justification must be borrowed from this or that philosophical theory.

Philosophy, on the other hand, encompasses activities in multiple contexts, and deals with objects for which multiple levels of description are necessary beyond the merely physical, material, or neurobiological. Philosophy requires reasoning, argumentation, evaluation, analysis and synthesis, interpretation, *hermeneutics*: it deals with broad concepts like *meaning, value, significance*. It is concerned with *social realities* and *abstract systems* no less than the biological, chemical and physical substrates of such higher-order structures. Philosophical accounts have to make sense of phenomena which are social, linguistic, cultural, political, experiential, aesthetic, evaluative, normative: these are generally not delineable and determinable the way the objects of the sciences are, and they are not explicable in the same terms. Scientific data may be relevant to such analyses; but they can never be decisive or dispositive: philosophy is larger than science, and *comprehends* it (in both senses of the word). The objects of philosophy include higher levels and larger contexts, the social structures, institutions, organized practices and collective functions—including *philosophy itself*—which happen not just in brains but in ‘the external world,’ with everything that entails.

In Heidegger’s terms, *Dasein* or ‘human being’ is always already being-in-the-world (*in-der-Welt-sein*), the phenomenology of which requires *philosophical* investigation and explanation: a full accounting and reflective interpretation which provides for the ‘lifeworld’ within which human being takes place, the multiple contexts larger than the given physical immediacies of

brain activity. Philosophy can, where neurophilosophy cannot, address the reality of complex, externalized, formal systems, structures, and narratives, transcending individuals and encompassing not just ‘many brains’¹—the fallacy of composition—but going beyond brain activity altogether.

Back to Kant, Hegel, and Heidegger

I have taken note of some clear parallels between contemporary antireductionist arguments and important positions maintained by Kant, Hegel, and Heidegger. A more extensive treatment of this topic would bring out the many detailed points of similarity connecting Nagel and Chalmers with their predecessors in the German-language philosophical tradition, who also include (among others) Leibniz, Schelling, Nietzsche, Husserl, and in his own way, Wittgenstein.

Here at least a few of these parallels may be briefly drawn. In Hegel’s analysis of *explaining consciousness*, for example—what may be his most essential philosophical project—a diremptive yet self-reflexive relation obtains, between the conscious *object* which is to be explained, and the conscious *subject* which does the explaining: the same thing accounting for itself to itself. We have to explain ourselves, as the knowing, experiencing beings who are capable of explaining and accounting. Making intelligible one’s own being and knowledge of oneself must be part of one’s general account of the world. We are the being that does the explaining of being, including our own: we are the irreducible consciousness that has come to understand itself as irreducible consciousness. Knowing ourselves for what we are, the self-conscious subjects as well as the objects of inquiry and knowledge, clearly suggests Hegel’s Mind or Spirit, coming to know itself as Mind or Spirit—as ‘substance just as much as subject.’ Nagel (2012) acknowledges his debt to Hegel and German philosophy, sounding especially Hegelian when he writes (for example):

Mind ... is doubly related to the natural order. Nature is such as to give rise to conscious beings with minds; and it is such as to be comprehensible to such beings. Ultimately, therefore, such beings should be comprehensible to themselves.

The subjectivity implicit in self-consciousness and self-comprehension requires the reality of *reflexively self-aware persons*, conscious *selves*, who have *first-person experience*. Following Kant we could even call these the *a priori* conditions for the possibility of explaining consciousness. Not only is the subject contained *a priori* in any and all such explanations; so too is the

¹Note Churchland’s (2008) deformation of the language of metaethics and social theory into terms of brain activity: ‘Solving social problems is an awesomely complex business, requiring relevant facts, including facts about cultural practices, about what brains do value, and fact-based predictions about consequences. Fundamentally, moral/social problems are *constraint-satisfaction problems* at the many-brain level...’

self-conscious rational *community of inquirers*, for which truth, knowledge, scientific explanation itself are all real values. Only on such an assumption does the entire explanatory enterprise make any sense. These are among the phenomena of consciousness that need to be explained, to which no reduction to brain activity could possibly be adequate.

German-language philosophers had long recognized and argued that empirical research into psychology, anthropology, biology and the natural sciences (necessary as it is) would never—*could* never—be sufficient for a philosophically adequate account and comprehensive explanation of ‘mind,’ self-consciousness, thought, and reason. But philosophy in the English-speaking world had to come around to these conclusions independently, that the subject of first-person experience is ‘transcendental’ in Kant’s sense—that self-consciousness is ontologically irreducible, that it is no epiphenomenon but fundamental to the nature of the universe. The results, as Chalmers (2010) writes, are ‘liberating’:

This I take to be precisely the liberating force of taking consciousness as fundamental. We no longer need to bash our heads against the wall trying to reduce consciousness to something it is not. Instead we can engage in the search for a constructive explanatory theory ... which accommodates consciousness in the natural world. And a fundamental theory of consciousness ... is the best way to do just that.

Assigning consciousness its proper place as fundamental, the consequences are profound: we have to expand our ontology and become realists about a broad range of possible objects, not all of which are physical. This will require a substantial revision to our entire conception of the universe: if consciousness is real, then materialism as the whole truth is ‘almost certainly false.’ Just as with Kant, Hegel, and Heidegger, reason, ontology, phenomenology, and social theory are all advanced again as central and necessary, as ineliminable in the project of explaining and accounting adequately for consciousness. These long-held ‘continental’ positions may now be said to have been vindicated from within the analytic philosophical tradition itself.

Bibliography

- Chalmers, D. (1996). *The conscious mind. In search of a fundamental theory*. New York: Oxford University Press.
- Chalmers, D. (1997). ‘Facing up to the problem of consciousness.’ In: J. Shear (ed.), *Explaining consciousness. The ‘hard problem,’* 9-30. Cambridge: MIT Press.
- Chalmers, D. (2010). *The character of consciousness*. New York: Oxford University Press.
- Churchland, P.M. (1988). *Matter and consciousness. Revised edition*. Cambridge: MIT Press.

- Churchland, P.S. (1989). *Neurophilosophy. Toward a unified science of the mind/brain*. Cambridge: MIT Press.
- Churchland, P.S. (2008). 'The Impact of Neuroscience on Philosophy.' *Neuron* 60 (Nov): 409-411.
- Dennett, D. (1986). *Content and consciousness* (2/e). London and New York: Routledge.
- Dennett, D. (1991). *Consciousness explained*. New York: Little, Brown.
- Nagel, T. (2012). *Mind and cosmos. Why the materialist neo-Darwinian conception of nature is almost certainly false*. New York: Oxford University Press.
- Shear, J., ed. (1997). *Explaining consciousness. The 'hard problem'* [articles on Chalmers from the *Journal of Consciousness Studies*]. Cambridge: MIT Press.