

Is Emergent Anomalous Panpsychism Viable?

(draft)*

David Bourget

February 2, 2017

Abstract

We can classify theories of consciousness along two dimensions. The first dimension is a theory's answer to the question of whether consciousness is "something over and above" the physical. Physicalism, dualism, and Russellian monism are three possible positions on this dimension. The second dimension is a theory's answer to the question of how conscious states causally interact with physical states. The three main possible answers to this question are nomism (the two interact through deterministic laws), acausalism (they do not causally interact), and anomalism (they interact but not through deterministic laws). This paper explores the potential and viability of anomalous dualism, a combination of views that has not been explored. I suggest

*Forthcoming in William Seager, ed., *The Routledge Handbook of Panpsychism*.

that a specific version of anomalous dualism, emergent anomalous panpsychism, can address the two most pressing issues for dualist views, the problem of mental causation and the mapping problem (the problem of predicting mind-body associations). Emergent anomalous panpsychism seems to be the only theory that can reconcile all the evidence that has been offered by dualists and physicalists.

To the extent that the mind-body problem is considered to be a problem, it is because it pits against each other two at least somewhat persuasive lines of argument. On the one hand, there are well-known arguments purporting to show that the conscious mind is something over and above physical going-ons, a view known as *dualism* about consciousness. On the other hand, dualism does not seem to square with the picture of the world painted by the biological and physical sciences. Aside perhaps for small amounts of indeterminacy (which is *prima facie* irrelevant), every physical event seems to be determined by prior physical events. This seems to leave no room for non-physical minds to play a causal role in the flow of physical events. But it seems that our minds do play a causal role in the flow of physical events. These observations seem to favor *physicalism* about consciousness, the view that the conscious mind is nothing over and above the physical.

For someone who is even just a little bit moved by the two sides of this debate, accepting one argument to the detriment of the other represents a compromise. We should consider compromises of this sort only when we have truly exhausted the options, and I am not convinced that we have. At

any rate, my aim in this paper is to explore the few options that are open to us if we don't want to compromise significantly on the premises of the preceding arguments or any widely accepted empirical evidence. This will involve consideration of some rather exotic alternatives to the standard views. The first alternative I will consider is Russellian constitutive panpsychism, which has attracted considerable attention in recent years. As some authors have pointed out already, this view is ruled out by the same considerations as physicalism. This will lead us to consider an even stranger view: anomalous dualism. Anomalism is the view that mind-body interactions are not governed by deterministic laws. The combination of anomalism and dualism seems to not have been explored. This is presumably because this view faces some obvious challenges: it seems to be inconsistent with such mundane observations as that our conscious states cause bodily movements and always occur in the same conditions. I will suggest that a specific kind of anomalous dualism, emergent anomalous panpsychism, is consistent with such observations and escapes the dilemma that the standard arguments against physicalism and dualism pose. I will also suggest that the hardest part of the hard problem for standard physicalist and dualist views is what I call the *mapping problem*, the problem of building a fully general predictive model of mind-body associations. Anomalous dualism has the virtue of dissolving this problem.

I begin by reviewing the dilemma that the standard arguments against physicalism and dualism pose, along with its failed Russellian resolution.

1 Physicalism, Dualism, and Russellian monism

In the introduction, I glossed physicalism about consciousness as the view that the conscious mind is nothing over and above the physical, and dualism as the view that the conscious mind is something over and above the physical. These are intuitive glosses on these views that are widely accepted, but for developing clear arguments it is necessary to work with sharper definitions. The following is a sharpening of the general physicalist thesis that most physicalists accept: the complete physical truth about the world necessitates all truths about the world. Physicalism about consciousness is then the view that the complete physical truth about the world necessitates all the phenomenal truths about the world. We can think of dualism about consciousness as the view that phenomenal truths are not necessitated by physical truths.

1.1 The challenge for physicalism

Physicalism has been targeted by a number of arguments that rest on the existence of an epistemic or inferential gap between mental and physical facts, in particular, the argument from disembodiment (Descartes 1641), the knowledge argument (Jackson 1982), and the conceivability argument (Kripke 1980, Chalmers 1996, 2003a). Here I am going to focus on Chalmers' elaboration of the conceivability argument.

Let P stand for a complete statement of all physical facts about the world. Let Q be the claim that there is consciousness. Then the statement

$P \& \neg Q$ describes a *zombie world*, a world physically like ours but without consciousness. Chalmers' argument goes:

1. $P \& \neg Q$ is conceivable
2. If $P \& \neg Q$ is conceivable, $P \& \neg Q$ is possible.
3. If $P \& \neg Q$ is possible, physicalism is false.
4. Therefore, physicalism is false.

Conceivable here means that the statement cannot be ruled out a priori (it is not a priori false). Premise 1 is fairly plausible: it does not seem that a priori reasoning alone can rule out the possibility of a zombie world. Premise 3 is true by definition on our understanding of physicalism, assuming that Q is true (that there is consciousness). Clearly, premise 2 is the one that bears most of the weight of the argument. Chalmers has two related arguments for premise 2. The first is an inductive argument for a general claim that entails premise 2 given physicalism:

Scrutability of Truth For any statements D and S , if D is a complete qualitative description of the actual world and S is true, then $D \supset S$ is a priori.

To a first approximation, a *qualitative* description is one that contains no names but only predicates, quantifiers, and indexicals. This qualification prevents *The actual world is @*, where @ is the name of the actual world,

from counting as a complete description of the world, which would clearly refute Scrutability. A *complete* qualitative description of the world is roughly one that necessitates all facts about the world. According to Chalmers (1996, 2002, 2003a) and Chalmers and Jackson (2001), Scrutability is confirmed by every uncontroversial truth, including in particular truths that are necessary but a posteriori. For example, $P \supset (water\ is\ H_2O)$ is a priori (even though *water is H₂O* by itself is not).

The intuitive idea behind Scrutability is that a truth is a posteriori when we lack sufficient empirical information to determine that it is true. Given sufficient empirical, qualitative information of the kind that we can acquire through our senses and experimentation, we should be able to tell what the extensions of our terms are, and this should allow us to know all truths. For example, once I know what are the extensions of “Dog” and “Fido”, I am in a position to tell whether or not Fido is a dog. This seems to be generally true, so Scrutability seems warranted.

Chalmers’ second line of argument for premise 2 appeals more directly to (one version of) his two-dimensional semantic framework. It can be put intuitively as follows. When we consider whether $P \mathcal{E} \neg Q$ is a priori, we grasp a certain way the world could be, and we reason on this in search of an inconsistency. We can think of what we grasp as a set of possible worlds or a function from possible worlds to truth values (an intension). We can call this the *epistemic intension* of the sentence. There is also a way the world could be that is relevant to the metaphysical possibility of $P \mathcal{E} \neg Q$. Again this is a

set of possible worlds or intension, which we can call its *subjunctive intension*. Now, if what we grasped when checking $P \mathcal{E} \neg Q$ for a priori inconsistency was the very same intension that is relevant to the metaphysical possibility of $P \mathcal{E} \neg Q$, then this statement not being a priori false would imply that it is not necessarily false. Chalmers argues that $P \mathcal{E} \neg Q$ has identical epistemic and subjunctive intensions. The reason is that we have a full, a priori grasp of the natures of the relevant physical and phenomenal properties.¹ This is not true of arbitrary statements, but it is true of $P \mathcal{E} \neg Q$ because of our special access to phenomenal and physical properties. (We will come back to this issue in 1.3). This means that the conceivability of $P \mathcal{E} \neg Q$ implies its possibility.

One objection to the conceivability argument points out that P is vastly complicated. It might even be in some sense as large as the world. *That's* a complicated sentence to parse. Moreover, since physics does not seem to be a finished project, it does not seem that we even know what sorts of properties would ultimately figure in a complete account of the microphysical. How, then, can we confidently assert that $P \mathcal{E} \neg Q$ cannot be ruled out a priori? We only have a tiny glimpse of what might be in P .

There are two parts to this worry: the first part is that we don't know what the final physics might look like; the second part is that we don't have the mental capacity to entertain claims anywhere near the complexity of P . The answer to the first worry is that, even if we don't know which of the

¹See Chalmers 2003b, 2007, 2009 for details.

metaphysically possible physical properties are instantiated, we can consider any such property hypothetically, and we can see that, if the final description of world attributed this property to certain things, such attribution would not a priori entail Q . Since this holds true of arbitrary physical properties, it seems that $P \not\mathcal{E} \neg Q$ will be conceivable whatever P turns out to contain. The second worry is addressed by considering increasingly complex combinations of properties. Since it seems that moving to more complex physical descriptions does not get us any closer to a description of the world that a priori entails Q , it seems reasonable to conclude that no combination of such properties will do.

1.2 The challenge for dualism

If physicalism is false, it seems, dualism must be true: phenomenal states are something over and above physical states. But there is a compelling case against dualism. First, we know that pretty much everything that happens in the physical world is nomologically determined by antecedent physical events. Famously, microscopic, quantum-level events are not always determined in this way, but this does not seem relevant. Something like the qualified Completeness claim below seems to be true at least.

Completeness Aside for partially random effects, all physical events are nomologically determined by prior physical events.

A physical event is an instantiation of a physical property (perhaps by an

individual at a time).

It is plausible that phenomenal events (i.e. instantiations of phenomenal properties) at least sometimes make a difference to the course of physical events. For example, it seems that I wouldn't have taken an Advil if I had not felt a headache. Thus Efficacy seems plausible.

Efficacy Phenomenal events sometimes make a difference to the course of physical events.²

I take making a difference to entail counterfactual dependence relations. If Efficacy is true, claims of the following form are true for some mental and physical events (my claim about Advil being a plausible example):

P-M-Dependence Had mental event M not occurred, physical event P would not have occurred.

Efficacy, Completeness, and dualism seem in tension, but they are not strictly speaking inconsistent because our Completeness principle contains an exception. To generate an inconsistency, we need an assumption to the effect that the exception does not apply:

Mind-to-Matter Nomism The effects of phenomenal events on physical events are determined by strict deterministic laws.

²Several authors formulate the problem of mental causation for dualism in terms of a No Overdetermination principle. In effect, I am by-passing such considerations by packaging a ban on systematic overdetermination as part of my Efficacy claim (Efficacy is inconsistent with all mental events being overdetermined, as overdetermination is normally understood).

Given Mind-to-Matter Nomism, we can derive a contradiction as follows. Efficacy entails that claims of the form of P-M-Dependence are true. From Mind-to-Matter Nomism, we know that the event P mentioned in P-M-Dependence was not a partially random effect that is excluded from the scope of Completeness. From this and Completeness, we know that P was determined by a prior physical event C. Now the key question is whether C would have occurred even if M had not occurred. If this is the case, then P-M-Dependence is false, because C is nomologically sufficient for P. In assessing counterfactuals such as P-M-Dependence, we must imagine a situation that is as much as possible like the actual world consistently with M not occurring in it, and we must check whether P occurs in this situation. So, unless C necessitates M, the M-less situation that we have to consider when assessing P-M-Dependence is one where C occurs. And we know from dualism that C, a physical event, does not necessitate any phenomenal event. So we know from dualism together with Completeness and Mind-to-Matter Nomism that, had M not occurred, P would have occurred anyway. This contradicts P-M-Dependence. So Efficacy, Completeness, Mind-to-Matter Nomism, and dualism are inconsistent.

This leaves us with a conundrum: on the one hand, the conceivability argument seems to rule out physicalism; on the other hand, dualism is inconsistent with plausible claims about mental causation. Russellian monism offers an apparent solution to this dilemma, which I discuss in the next section.

1.3 Russellian monism and its conceivability challenge

The Russellian monist’s purported solution to the preceding dilemma turns on a distinction between two different kinds of physical properties (or two kinds of properties that one might reasonably take to deserve the label “physical” as it is generally intended to be used). *Narrowly physical properties* are properties of roughly the same kind or flavor as have so far been uncovered by biology, chemistry, and physics (to a first approximation, dynamical-structural properties). For my purposes here, I am going to count higher-order or functional properties realized by narrowly physical properties as also narrowly physical. *Broadly physical properties* include both narrowly physical properties and the properties that actually realize narrowly physical properties (if any).³ Some properties might be broadly physical but not narrowly physical. If, as Russell (1927) and Maxwell (1971 and 1979) suggest, the fundamental properties revealed by physical science are mere dispositional properties, it might be that there is an underlying layer of categorical properties that realize these properties. These hidden realizers, if they were different in kind or flavor from the properties revealed by physical science, would not fall in the category of narrowly physical properties, but they would be broadly physical.

Given these two possible understandings of physical properties, we have two possible understandings of physicalism. *Narrow physicalism* holds that phenomenal properties are determined by narrowly physical properties. *Broad*

³The distinction between narrowly and broadly physical properties is made in those terms by Chalmers (2015), but Stoljar (2001) puts forward a closely related distinction.

physicalism holds that they are determined by broadly physical properties. One can be a broad physicalist while rejecting narrow physicalism (but not vice versa). This is *Russellian monism*. Thus, the distinction between narrowly and broadly physical properties leaves us with three possible views: narrow physicalism, Russellian monism, and dualism understood as the view that mental properties are not necessitated by broadly physical properties. (I have been using “physicalism” (unqualified) to mean narrow physicalism, and I will continue to do so.)

Several authors have remarked that Russellian monism, a kind of broad physicalism, seems to escape the conceivability argument.⁴ It escapes the argument as stated above if we take P to be the totality of *narrowly* physical truths about the actual world. Since it is part of the Russellian monist view that there is more to the physical than the narrowly physical truths capture, this view leads us to deny premise 3: (broad) physicalism might be true consistently with the possibility of narrow zombies.

Russellian monism also seems to avoid the argument from Completeness. Completeness and Efficacy are true on either reading of “physical”. However, the Russellian monist can plausibly say that, had M not occurred, C would not have occurred. This is because the categorical basis of C might well necessitate M , and if C is a mere disposition as the Russellian monist claims, it cannot exist without a categorical basis. Compare with a clearly dispositional

⁴Chalmers (1996, 2003a, 2015) explicitly makes this point. Strawson’s (1994, 2003 and 2006) case for panpsychism effectively turns on this point.

property such as fragility. In the case of fragility, the following counterfactual seems trivially true: had all the fragile objects not existed, nothing would be fragile. A counterfactual world with no fragile objects is a world without fragility, not a world with fragility floating free from its categorical base. Similarly, if the Russellian monist is right about the nature of C and the relationship between its categorical base and M , it seems that we should think of the nearest M -less world as a C -less world.⁵

We can distinguish two brands of Russellian monism: *constitutive Russellian panpsychism* (or just *Russellian panpsychism*) and *panprotopsychism*. The former takes the hidden inner natures of microphysical properties (or at least some such hidden natures) to be phenomenal, whereas the latter does not. On the latter view, narrowly physical properties and their inner natures together necessitate consciousness as we know it, but neither are themselves phenomenal nor narrowly physical.⁶

I am going to set panprotopsychism aside for the rest of this paper. My main reason for doing so is that this view seems to offer no explanation at all. We have not been given a theory of consciousness until we are told what lies at the heart of microphysical properties such that they can generate consciousness as we know it. Until we are told *what* it is and *how* the unknown

⁵However, one might think that narrowly physical properties are properties of another kind that can be instantiated without further grounds. This would threaten to make phenomenal properties epiphenomenal on the Russellian monist view. Here I am going with what seems to be the least problematic Russellian view.

⁶Another possible view is that they are narrowly physical, but this seems to be no improvement over narrow physicalism.

ingredient does the job, and not merely *that* it does the job, we have not been given an alternative theory of consciousness. Panprotopsychism is no more illuminating than *otherism*, the view that what generates consciousness as we know it is something other than what we have considered so far. In practice, our options are Russellian panpsychism, dualism, and narrow physicalism.

Russellian panpsychism seems promising at first, but it arguably succumbs to a revised conceivability argument. The conceivability argument against panpsychism arises as part of a broader challenge known as *the combination problem* (Seager 1995, 2002). It seems exceedingly implausible that the fundamental particles bear *our* phenomenal properties, for example, the property of experiencing a computer screen with such and such shapes on it. More plausibly, the microphenomenal properties are very simple and quite unlike the phenomenal properties with which we are familiar. This means that the panpsychist has to explain how macrophenomenal properties of the kind we are familiar with arise from microphenomenal properties together with other broadly microphysical properties. This is the combination problem.

We can distinguish several aspects of macrophenomenal experiences that are on the face of it non-trivial to explain on the panpsychist view. First, it seems plausible that macrophenomenal properties are more complex than microphenomenal properties: they have a kind of internal structure that is plausibly not there at the microphenomenal level. For example, I am experiencing the screen *on top* of the desk. The panpsychist needs to account for how relatively simple microphysical and microphenomenal properties can

together constitute complex phenomenal properties. It also seems plausible that we experience qualities that are not experienced by fundamental particles. For example, it seems implausible that electrons can have stabbing pains. So macrophenomenal experiences not only involve more complexity, but also new qualities. It also seems that we are subjects of experience above and beyond fundamental particles. Again, we need to explain how these macrosubjects arise.⁷

The claim made by Russellian panpsychism that distinguishes it from dualism is that all consciousness is necessitated by the conjunction of the narrowly microphysical truths and the microphenomenal truths. It is easy to see that the combination problem can be turned into a conceivability argument against this claim. This point is developed at length by Goff (2009), so I am going to go over it somewhat quickly.

Take the totality of microphenomenal and narrowly microphysical truths about the world, $P+$. As we just observed, it seems overwhelmingly plausible that macrophenomenal properties are distinct from microphenomenal properties. So there are macrophenomenal truths that are not included in $P+$. The question is whether they are metaphysically necessitated by $P+$ as asserted by Russellian panpsychism. Take any macrophenomenal truth Q . A conceivability argument that parallels the argument against physicalism can be made as follows:

1. $P+ \& \neg Q$ is conceivable

⁷This way of charting the problem comes (roughly) from Chalmers 2015.

2. If $P+ \mathcal{E} \neg Q$ is conceivable, $P+ \mathcal{E} \neg Q$ is possible.
3. If $P+ \mathcal{E} \neg Q$ is possible, then Russellian panpsychism is false.
4. Therefore, Russellian panpsychism is false.

Premise 1 is *prima facie* plausible. As Goff argues, it seems easy to imagine a microexperiential zombie world, a world where basic particles such as electrons have simple experiences, but there are no macroscopic consciousnesses with complex human experiences. There seems to be nothing in the idea of tiny particles having experiences and being organized in the sorts of ways that we know particles are organized in the body that requires us to think of these arrangements as having human-like consciousnesses. Premise 3 is true by definition of panpsychism (assuming that Q is true). As in the case of physicalism, premise 2 is where most of the weight of the argument lies.

Premise 2 seems to be justifiable in all the same ways as the parallel claim regarding $P \mathcal{E} \neg Q$. In particular, the independently justified Scrutability principle applies just the same if we take D to be $P+$. Since panpsychism claims that $P+$ entails all phenomenal truths and we know independently that non-phenomenal truths are entailed by microphysical truths, we know that $P+$ entails all truths if panpsychism is true. This allows us to apply Scrutability as in the case of physicalism.

It seems that the only way the panpsychist might escape this argument while continuing to maintain (in a weakened form) the anti-physicalist argument that motivates her position is by rejecting Scrutability and arguing that

$P+$ is not sufficiently grasped a priori for the conceivability of $P+\mathcal{E}\rightarrow Q$ to entail its possibility. To put things in descriptivist terms, this would mean that $P+$ involves terms that pick out their referents by contingent features.

In the case of physicalism, this kind of concern arises naturally from the fact that fundamental physics does not seem to describe the inner nature of fundamental physical properties. This introduces the possibility that there is more to the world than the narrowly physical facts because there are underlying natures that are not captured by such facts. This is the panpsychist's interpretation of the upshot of the argument against physicalism. But the panpsychist purports to tell us what are these underlying natures. Once this has been done, we can no longer say that there is some unknown stuff that plays a role in determining macrophenomenal properties. If the panpsychist admitted that there is further layer of reality that needs to be specified in order to capture all aspects of reality relevant to explaining the macrophenomenal, this would be an admission that the narrowly microphysical and the microphenomenal do not necessitate the macrophenomenal.

Such considerations might lead one to revisit premise 1. One might point out that the complete microphenomenal and microphysical description of the world would be utterly complex, and that we have little clue what microphenomenal properties might figure in it. These points might be taken to suggest that, despite appearances, $P+\mathcal{E}\rightarrow Q$ might not be conceivable to a fully rational being even if *we* can't find any inconsistency in it.

The preceding points parallel similar remarks that we considered in the

case of the argument against physicalism. Similar replies come to mind. First, we might not know what specific microphenomenal properties there are (just like we don't know what the final physics will be like), but any example that we can consider hypothetically fails to a priori entail macrophenomenal facts. For example, little charge buzzes just don't add up to a visual scene. There seems to be no reason to think that more variations on the same theme would help. Regarding the complexity of $P+$, here too it seems that considering increasingly complex microphenomenal and microphysical facts does not take us any closer to an explanation of the macrophenomenal. This suggests that $P+\mathcal{E}\rightarrow Q$ really is conceivable.

1.4 The nomism question

All positive views seem to be ruled out: physicalism and Russellian panpsychism are ruled out by conceivability arguments, and dualism is ruled out by the argument from Completeness. At this point there seems to be no more distinctions that one can make to introduce new options that would avoid the horns of the dilemma. Since the three arguments discussed in the three preceding sections are valid, at least one of the premises of these arguments has to go.

Many philosophers who have thought about these matters have made up their mind and decided which premise they deny. Today, the most commonly endorsed position is physicalism (broad or narrow) combined with the denial of premise 2 of the conceivability argument against either broad or narrow

physicalism. Smaller numbers of broad or narrow physicalists deny other premises of the conceivability arguments. There are two dualist positions with a following: dualism without Completeness (standard interactionism) and dualism without Efficacy (which can be either a form of epiphenomenalism or a form of parallelism).

Each of the preceding positions has a considerable associated cost, because the premise that each denies has considerable evidence going for it. Premise 2 of the conceivability argument might be the most tempting to deny because the evidence for it is largely a priori and philosophical in nature, but it is compelling evidence nonetheless. Evidence for such quasi-logical truths is always a priori or philosophical in nature. It would be a mistake to think that premise 2 of the conceivability argument (or any other premise) has little going for it because many deny it. The fact that many deny it at best suggests that it is the *least* well supported premise, not that it is not well supported. As an analogy, consider that when two vehicles collide head-on, there is always one that gets pushed back, but this doesn't mean that it is a light vehicle.

Completeness and Efficacy also seem costly to deny. Efficacy seems to be strongly supported by mundane observations as well as general theoretical considerations. Among other theoretical considerations, it seems that if consciousness made no difference, we shouldn't expect brains created through natural selection to involve any consciousness. Such considerations make denying Efficacy seem unappealing. Denying Completeness seems to fly in

the face of very well confirmed scientific theories.

There is one premise of our conundrum that has not been closely scrutinized: Mind-to-Matter Nomism. A dualist who denied this claim would be in a position to resist the argument from Completeness at least as formulated here. Of course, Mind-to-Matter Nonism is not implausible on its face, but it has not been discussed very much, so one might hope that its apparent plausibility will dissipate upon closer inspection.

Philosophical discussions of consciousness tend to focus on the choice between dualism, physicalism, and Russellian panpsychism, but the role of Mind-to-Matter Nonism in the argument from Completeness against dualism draws our attention to another dimension of variation for theories of consciousness: it could be that phenomenal properties causally interact with physical properties through deterministic laws of nature, it could be that phenomenal properties don't causally interact with physical properties, or it could be that phenomenal properties causally interact with physical properties but not through deterministic laws. I am going to refer to the first view as *nomism*, the second as *acausalism*, and the third as *anomalism*. One can be a nomist (or an anomalist or an acausalist) about some interactions and not others, but, presumably, mind-to-matter interactions are either generally nomic or generally anomalous, and matter-to-mind interactions are either generally nomic or generally anomalous.

While the choice between physicalism and dualism pertains to the existence and types of necessitation relations between mental and physical

properties, the choice between nomism, anomalism, and acausalism pertains to the existence and types of causal relations between mental and physical properties. Prima facie, these dimensions of variation are independent of each other. In practice, however, most physicalists (broad or narrow) seem to be nomists in both directions. Dualists hold and discuss a wider variety of views. As mentioned above, three views have traditionally had some currency: interactionism, epiphenomenalism, and parallelism. Interactionism (at least in its standard form) endorses nomism about both mind-to-matter and matter-to-mind interactions, positing psychophysical laws that govern these interactions. Epiphenomenalism endorses nomism about matter-to-mind interactions while denying that there are mind-to-matter interactions. Parallelism is acausalist in both directions. The view that has not been considered so far is dualism combined with anomalism about mind-to-matter interactions (or both directions). Since it is (as far as I can tell) completely unexplored, it is natural to ask whether this view has any potential to yield a satisfactory theory of consciousness. I will consider this question in section 3. Before doing so, I will explore a widely neglected aspect of the mind-body problem that invites us to reconsider nomism in the matter-to-mind direction.

2 The mapping problem

Physicalism, Russellian monism, and nomic dualist views have an important implication in common: they imply that there should be general, at least

nomically necessary principles predicting which phenomenal properties occur in which physical conditions. On reductive physicalist views, the case of consciousness is supposed to be like the case of heat, water, life, or anything else that has been successfully reduced to physical properties: we expect some kind of general translation of consciousness talk to physical talk to be possible in principle. If such a translation were possible, we could formulate general principles pairing phenomenal properties (i.e. physical properties under a phenomenal description) with physical properties (i.e. phenomenal properties under a physical description). On non-reductive physicalist views, translatability between mental language and the language of physics is not expected, but translatability between mental language and a more abstract functional language is expected, so we should at least be able to spell out principles that connect mental states with physical states broadly construed to include functional states. Whether physicalism or a dualist view that endorses nomism about matter-to-consciousness interactions is correct, we can expect to eventually find general principles allowing us to predict which phenomenal descriptions apply to which physical descriptions across all physically possible circumstances. I will refer to such general principles as *psychophysical mappings*. The problem of specifying a psychophysical mapping is the *mapping problem*.

The mapping problem is distinct from, and plausibly largely independent of, the question that opposes physicalists and dualists. First, it is clear that knowing the correct psychophysical mapping would not automatically tell

us what is the correct metaphysics of consciousness, though it might suggest either physicalism or dualism.⁸ Conversely, neither physicalism nor dualism gives us or even suggests a psychophysical mapping. Consider reductive physicalism, the claim that phenomenal properties are physical properties. This view tells us that there should be a psychophysical mapping, but it does not tell us what the mapping is, because it does not tell us to which physical property a given phenomenal property is identical. In this respect, the claim that phenomenal properties are physical properties differs from the claim that heat is identical to mean kinetic energy: in the latter claim, there is an implicit specification of which instances of mean kinetic energy any given instance of heat is identical to (the more heat, the more kinetic energy), but there is no parallel mapping implicit in the claim that phenomenal properties are physical properties.⁹ The latter claim leaves us entirely in the dark regarding the correct psychophysical mapping. This brings into focus the fact that it is possible to settle the debate between physicalists and dualists (say, by making a good case against dualism) without settling the mapping problem. Relatedly, the mapping problem seems to be the really hard part of the hard problem: a majority of philosophers of mind seem to think that they know that physicalism is true, so they think they have a solution to the traditional

⁸For example, if the correct mapping was generated by a principle that, in its simplest statement, covers all metaphysically possible cases, this might suggest that physicalism is correct. The correct mapping might also be obviously non-generalizable to physically impossible cases, which might seem to suggest that dualism is true.

⁹I am ignoring the fact that two magnitudes that range from infinitely small to infinitely large can be mapped 1-to-1 in infinitely many ways consistently with this constraint.

mind-body problem, but they don't have a clue how phenomenal properties map onto physical properties.

This last point should not be conflated with Levine's (1983) point that, even if knew that physicalism is true, we might will want to ask *why* such and such mental state is identical to such and such physical state. The mapping problem is not a problem about explanation but about prediction. It is conceivable that we might find a psychophysical mapping that successfully predicts all mind-body associations but that would not answer Levine's question; the mapping might still seem to call for an explanation. Conversely, it is conceivable that we might be able to answer questions such as Levine's (why are this phenomenal state and this physical state identical?) without being able to produce a psychophysical mapping.¹⁰ Like the question pertaining to the metaphysical status of consciousness, the explanatory gap seems to be independent of the mapping problem.

Neuroscience has revealed numerous correlations between brain areas and types of conscious experience and other kinds of mental activity. It has also revealed what appear to be limited mappings between aspects of conscious experience and certain kinds of brain activity, for example, the phenomenological color space can plausibly be mapped in a straightforward way to dimensions of activation in certain neural networks (see Churchland 1986). These are impressive findings, but they fall far short of a complete psychophysical mapping. By and large, the associations we known exist

¹⁰The story I sketch in the next section partly illustrates this possibility.

between phenomenal properties and physical properties don't seem to fall under a broad pattern that would allow us to specify a general psychophysical mapping. If we were to plot the known correlations between physical and phenomenal properties, we would see some local patterns (as in the case of color experience), but, aside from these patterns, the points would jump all over the place, forming no recognizable curve that we can characterize. We simply don't know how to extrapolate a general psychophysical mapping from what we have. Such a mapping has not even been *imagined*.

While physicalism *per se* does not address the mapping problem, there are some physicalist theories that might seem to specify a psychophysical mapping. All such theories fall under the broad heading of *tracking representationalism*. On these views, every phenomenal property is identical to a broadly physical property of the form *being in an internal state that has functional-physical feature F and tracks external physical state of affairs S*. The tracking relation that figures in such properties is supposed to be cashed out in broadly physical terms. It is supposed to correspond to a reductive account of representation, and theorists differ on which is the correct account. Feature F, which is also specified differently by different theories (e.g. 60 hz oscillations, being part of the global workspace, being poised to impact cognition in the right way, etc.) is supposed to distinguish conscious from unconscious representation.¹¹

On the tracking view, each phenomenal property is a matter of tracking

¹¹Tracking representationalist views include those of Baars (1988), Crick and Koch (1998), Dretske (1995), Lycan (1996), Rey (1995) and Tye (1995), among others. See Bourget and Mendelovici 2014 for an overview.

a specific external state of affairs while bearing feature F, for example, experiencing red is a matter of tracking things being red while having feature F, experiencing squareness is a matter of tracking squareness while having feature F, and so on. This might seem to specify a simple psychophysical mapping along the lines of this schema:

Representationalist mapping An individual experiences content C just in case they are in an internal state that has functional property F and stands in tracking relation R to C.

R is a placeholder for a specification of the relevant tracking relation, and F is a placeholder for the feature that distinguishes conscious from non-conscious representations. Contents might be entities of many different types (propositions, property complexes, properties of individuals, properties of worlds, etc.), so long as they can in some sense be tracked. Assuming that all experiencing can be understood as standing in some generic experiencing relation to a content-like entity, the representationalist mapping, together with a specification of R and F, seems to tell us in general what phenomenal properties occur under what physical conditions, just like the principle that heat is mean molecular kinetic energy tells us how much heat there is in any physical condition.¹²

The fact that every serious attempt at specifying a psychophysical mapping

¹²Some representationalists suggest that a further ingredient might be required: an intentional mode or representational manner (c.f. Crane 2003, Chalmers 2004, Speaks 2010, 2015). I argue against such extra ingredients in Bourget 2015, forthcoming a, and forthcoming b.

is an instance of the above schema makes tracking representationalism initially very attractive. The reason this approach seems promising is that, unlike mind-to-brain correlations, mind-to-environment correlations seem to fall under broad patterns: it seems that, at least in good conditions, our consciousness is filled with the very things in front of us. This works better for perceptual states than for non-perceptual states, but solving the mapping problem for perceptual consciousness would already be huge progress.¹³

Tracking representationalism is compelling at first, but it ultimately fails as a solution to the mapping problem. It faces two major issues. One issue is that the redness that enters our consciousness (what an experience of red is *of*) is not a physical property or any other property we might plausibly be said to track. The properties that are out there and available for tracking are properties such as *being disposed to reflect electromagnetic radiation of 650 nanometers*. We don't seem to be aware of electromagnetic properties such as this when we experience red, so, *prima facie*, experiencing red is not a matter of tracking such properties. Of course, this is similar to the conceivability argument against physicalism, and one might hope that similar replies will be available, but no one can deny the intuitive force of these considerations, and here I am interested in exploring the options that don't require us to reject such intuitions.¹⁴ In any case, the second issue for tracking

¹³Or one might be tempted to deny that there is non-perceptual consciousness, which many tracking representationalists have done.

¹⁴This argument is defended at length by Mendelovici forthcoming. Mendelovici & Bourget forthcoming presents this argument as well and argues that it is harder to block than arguments for dualism.

representationalism is more important to us here.

The second issue is that, even if tracking representationalism were true, it would not specify a psychophysical mapping. The kind of psychophysical mapping that ought to be available on a physicalist view is a mapping between phenomenal properties *under a phenomenal description* and phenomenal properties *under their physical description*. Tracking representationalism does not specify such a mapping. What it gives us is a mapping between physical descriptions such as “tracking Q while instantiating F” and partly physical, partly phenomenal descriptions such as “experience of Q”. Since Q is couched in physical terms (e.g. in terms of the language of electromagnetism) and the proper phenomenal description for the relevant experience is not (e.g. the proper description might be “experience of red”), this does not satisfy the requirements for a psychophysical mapping. This fact is easily overlooked because we happen to know, independently of tracking representationalism, that the experiences we have when looking at objects that reflect electromagnetic radiation of about 650 nm are experiences of red, but this is knowledge we have above and beyond what tracking representationalism tells us. By itself, tracking representationalism specifies no psychophysical mapping at all.

Once the apparent solution to the tracking problem offered by tracking representationalism is set aside, the mapping problem seems completely hopeless. There is simply no discernible general pattern in the known phenomenal-physical correlations revealed by neuroscience or everyday observation.

It is noteworthy that one aspect of the combination problem for Rus-

sellian panpsychism is a special case of the mapping problem. Russellian panpsychism, like other kinds of physicalism broadly construed, implies that there is a mapping between broadly physical properties and phenomenal properties. Part of the combination problem is that there seems to be no plausible way of specifying under what circumstances microphenomenal and narrow microphysical facts give rise to certain macrophenomenal facts. This part of the combination problem is a special case of the mapping problem we get when we assume that the items on one side of the mapping are combinations of microphenomenal and microphysical facts, as opposed to bare microphysical facts. This is a distinct problem from the other aspect of the combination problem mentioned earlier, which is a special case of the conceivability argument. In the same way that the general mapping problem might persist in the face of a resolution of the debate between physicalists and dualists, the mapping aspect of the combination problem might persist in the face of a solution to the part of the combination problem that arises from conceivability considerations, and it seems to be the hardest part of the combination problem.

3 Anomalous dualism

The mapping problem is a problem for views that imply that there should be a psychophysical mapping, i.e. physicalist views and dualist views that endorse Matter-to-Mind Nomism. The problem does not arise on a view that

rejects both physicalism and Matter-to-Mind Nomism. At the same time, the dilemma discussed in section 1 requires us to deny one of the premises of the arguments from conceivability and Completeness, and we have identified Mind-to-Matter Nomism as the least well established of these premises. When this premise is denied, dualism escapes the argument from Completeness. This naturally leads us to take a look at anomalous dualism, the kind of dualism that endorses anomalism about both mind-to-matter and mind-to-matter connections. To my knowledge, this view has never been considered before.

Anomalous dualism, like other forms of dualism, says that phenomenal properties are something over and above (narrowly or broadly) physical properties. However, it is immune to the argument from Completeness because it denies Mind-to-Matter Nomism. Anomalous dualism is also immune to the mapping problem. Since the anomalous dualist explicitly denies that there are psychophysical mappings, she does not have to worry about the apparent impossibility of finding a plausible mapping. On the contrary, the apparent hopelessness of this task is evidence for her view.

This is fine as far as it goes, but several objections spring to mind. First, by denying Mind-to-Matter Nomism, is not one accepting that consciousness is causally inefficacious in a different way (by having merely random effects)? This seems inconsistent with Efficacy. Second, how can the anomalous dualist explain the numerous correlations there are between mental states and physical states if there are no psychophysical laws or necessary associations? Third, if mind and body were only randomly associated, wouldn't our experience

be a mere “blooming, buzzing confusion”? Fourth, anomalous dualism might not have the problem of specifying a psychophysical mapping, but is this not simply because it gives up on explaining consciousness altogether? Lastly, one might doubt that this view is even compatible with Completeness as normally understood. After all, if the mind randomly affected the body, we would expect to witness macroscopic random events, which we never see. In sum, anomalous dualism is *prima facie* incompatible with Efficacy, Completeness, Mind-Brain Correlations, and Phenomenal Coherence, while at the same time seeming to give up on an explanation of consciousness.

Mind-Brain Correlations Phenomenal states and physical states are reliably correlated.

Phenomenal Coherence Our experience is generally coherent.

Despite appearances, anomalous dualism is at least in principle consistent with the preceding facts. The following made up story can serve as a kind of proof of concept, though its details are obviously implausible. Suppose that every physical property was “linked” to a randomly selected phenomenal property at the time of the Big Bang. Thereafter, linked phenomenal and physical properties have been associated and co-occurring across the universe. We can think of this initial random association as a fundamental principle of nature, a kind of stochastic law. Suppose also that the phenomenal properties of a physical system can in some circumstances have a random effect on the dynamics of the system. Say, for example, that any physical system about

to enter a total physical state involving physical properties associated with inconsistent phenomenal properties randomly jumps to another physical state, perhaps with probabilities determined by the physical state of the system.¹⁵ If the aforementioned random association had taken place at the time of the Big Bang and random jumps between physical states had been caused by inconsistent phenomenal properties since then, what should we expect to find today?

The story is underspecified in many respects, but it does not seem implausible that the principles it describes could, if precisified in the right way, result in a state of affairs that satisfies Phenomenal Coherence, Mind-Brain Correlations, Efficacy, and Completeness. The satisfaction of Mind-Brain Correlations is guaranteed simply by the initial stochastic association, which is consistent with anomalous dualism.

Phenomenal Coherence and Completeness require us to think about what would have happened over time if our story were true. Let us assume that there is some principled way of demarcating the relevant physical systems that counts properly functioning animal brains, or at least big parts of animal brains, as whole systems.¹⁶ Under this assumption, we would expect the

¹⁵We can say that two phenomenal properties are inconsistent when they are representations of inconsistent contents. This assumes that experiences are representational in some weak, non-reductive sense defended by Byrne (2001), Chalmers (2004), Pautz (2009), and myself (2010a, 2010b and 2015, forthcoming a, forthcoming b).

¹⁶Someone attracted to this idea might speculate that quantum entanglement is what delineates systems. A number of authors have explored entanglement-based explanations of the fact that consciousness seems to unify contributions from different parts of a physical system (e.g. Lockwood 1989, Penrose 1994, Seager 1995). There is a widespread misconception that “decoherence” virtually eliminates entanglement from the macroscopic world, but

brains of organisms to have evolved so that they barely ever enter physical states that correspond to inconsistent phenomenal properties (or at least not macroscopic states corresponding to inconsistent properties). Unless they had so evolved, they would be constantly suffering from the disruptive effects of random jumps, which would make them unstable, hence not prone to survive. Regarding simpler physical systems (including, in the limit case, isolated particles), their having fewer phenomenal properties (in virtue of having fewer physical properties) might explain their stability: they are not very likely to enter inconsistent states. The final result, then, should be brains and other macroscopic physical systems in which the potential random effects due to consciousness are largely absent, which is what we find. We should also expect the resulting stream of experiences supported by human brains to be generally coherent, which is what we find (Phenomenal Coherence). This picture respects Completeness as stated because it does not require consciousness to have any non-random effects, and it respects any similar principle that makes room for quantum indeterminacy because it does not require the random effects of consciousness to be more noticeable than the random effects predicted by quantum mechanics.

This story also predicts that phenomenal-physical associations should appear essentially random except perhaps for any structure implied by the

in fact the theory only predicts that decoherence makes entanglement unnoticeable by making the results of quantum measurement statistically like those of classical measurements. See Schlosshauer 2005 for a relatively non-technical explanation of what decoherence does and does not do.

original stochastic association principle. Again, this is roughly what we find. This last point is particularly interesting because no extant theory or sketch of a theory begins to explain the apparent arbitrariness of mind-body associations. Our made up story seems to provide not only a proof of concept for an explanation of the ways in which consciousness is organized, but it seems that a story along these lines could also potentially explain the ways in which it is *dis*organized. This might potentially answer the charge that such a theory cannot explain anything.

These observations are intriguing, but we have not yet shown that anomalous dualism is consistent with Efficacy. For anomalous dualism to be consistent with Efficacy as I understand it here, it would have to be consistent with the truth of claims of the form of P-M-Dependence, for example, Chocolate:

Chocolate Had you not been consciously thinking that there was a chocolate bar in front of you, you would not have reached out in the way you did.

It is possible to fill in the details of our story in a way that illustrates how Chocolate can be true consistently with anomalous dualism. Chocolate would be true if the story sketched above were true and two other conditions were met: a) your other phenomenal states (aside from your conscious thought that there was a chocolate bar in front of you) were inconsistent with all conscious thoughts whose contents have the form *there is an X in front of me*, where X is not a chocolate bar, and b) your brain was wired to produce a conscious thought whose content has the form *there is an X in front of me* if you are not in such a state. Condition (a) might obtain if, for example, your other

phenomenal states had the following contents: *I will touch X by extending my arm if there is an X in front of me; I will not touch any non-chocolate-bar by extending my arm.* Together, these contents are inconsistent with thoughts of the form *there is an X in front of me* where X is not a chocolate bar. Given (b), the antecedent of Chocolate (*had you not been consciously thinking that there was chocolate in front of you*) implies that you would have had some other conscious thought to the effect that something is in front of you.¹⁷ But given (a), such a thought would have been inconsistent with your other phenomenal states. The principles we are imagining predict that this would have resulted in disruption, and we could flesh out the story in such a way that this would have prevented you from reaching at just this moment. Of course, this example is highly contrived (I'm only trying to make a point of principle), but for all I know it could be that the brain has a massively redundant architecture that ensures that many phenomenal states cannot be altered without generating an inconsistency, which would underpin causal efficacy for these states on our story.

In sum, it seems that anomalous dualism, at least in the form specified here as an example (a kind of emergent anomalous panpsychism), can in principle

¹⁷Note that this does not require a “backtracking” evaluation of the counterfactual. Point (b) stipulates that your brain is wired to bring about a state of the relevant kind. Remove the conscious thought that there is a chocolate bar in front of you without backtracking, and from that point on the mechanism that generates a thought of the relevant form kicks in, generating it slightly later than the removed thought. That said, my intuition, for what it's worth, is that, depending on the circumstances, the nearest world where you don't think *there is a chocolate bar in front of me* could easily be one where you think that something else is in front of you.

explain surprising facts about consciousness and be made consistent with Efficacy, Completeness, Mind-Brain-Correlations, and Phenomenal Coherence. This takes away some of the initial implausibility of the view. Of course, the preceding discussion at best shows that the view is *in principle* consistent with what we know. The so so story that I have told is rather vague and implausible, and far from a complete theory of consciousness. More importantly, it is unclear whether our example emergent anomalous panpsychism or any similar theory can be squared with the (currently largely unknown) physical facts about the brain. The random effects that are left open by Completeness are supposed to be products of the measurement process in quantum mechanics, so a plausible development of anomalous dualism would have to align the random effects it posits with the predictions of physical theory, making sure in particular to respect the probabilities of random events that can be derived from the physics.

I don't know that this can be done, but it seems to me that there is no reason to be pessimistic at this time. In particular, anomalous dualism does not face the obvious objections from physics that extant "quantum theories of consciousness" face. Consider, for example, the theories of Eccles (1978) and Stapp (1996). These theories posit deterministic, principled effects from the mental to the physical. On Eccles' view, a conscious intention to move one's arm might cause neural firings that move one's arm by selecting which outcome of a quantum measurement process occurs. On Stapp's view, conscious intentions determine what quantum measurements occur in the

brain, which can guide its dynamic evolution. Here too the physical effects are supposed to correspond in a principled way to the mental states, for example, an intention to move one's arm corresponds to measurements that have to do with arms. These attempts to "sneak" deterministic mental-to-physical effects in the causal gaps left open by quantum mechanics are problematic because the gaps are just not big enough (Bourget 2004). In contrast, anomalous dualism does not try to sneak deterministic, principled effects among quantum effects. On a theory such as our example emergent anomalous panpsychism, there is no need for there to be any effects that make sense between the mental and the physical. Furthermore, the required effects are very slight compared to those required by a theory such as Eccles and Stapp. As we saw above, anomalous dualism can in principle explain all the salient facts about mind-body associations consistently with there being virtually no macroscopic random events where consciousness in fact makes a difference: consciousness only makes a difference in virtue of counterfactuals whose antecedents are generally false. As far as I know, there is no principled difficulty for anomalous dualism stemming from physics. Of course, it is also unclear whether the details can be worked out satisfactorily.

I tentatively conclude that there is no clear, commonsensical, or philosophical case against anomalous dualism: setting aside the details of the physics, which might be very complicated to work out either in favor or against the view, anomalous dualism accords with all the relevant considerations on hand, including conceivability considerations, Efficacy, Mind-Brain-Correlations,

Completeness, and Phenomenal Coherence. Being consistent with all considerations available to philosophers is not in general an impressive feature of a theory, but here it is noteworthy because no other view is. As we have seen, other views force us to reject conceivability considerations, Efficacy, or Completeness. Anomalous dualism also promises to offer an elegant explanation of the peculiar fact that mind-body associations seem largely random (if the view can be refined in a physically plausible way). These are remarkable and unique features of anomalous dualism that suggest that it deserves more investigation.

References

- Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.
- Bourget, D. (2004). Quantum leaps in philosophy of mind. *Journal of Consciousness Studies*, 11(12):17–42.
- Bourget, D. (2010a). Consciousness is underived intentionality. *Noûs*, 44(1):32–58.
- Bourget, D. (2010b). *The Representational Theory of Consciousness*. PhD thesis, Australian National University.
- Bourget, D. (2015). Representationalism, perceptual distortion and the limits of phenomenal concepts. *Canadian Journal of Philosophy*, 45(1):16–36.
- Bourget, D. (forthcoming a). Why are some experience "vivid" and others "faint"? representationalism, imagery, and cognitive phenomenology. *Australasian Journal of Philosophy*.

- Bourget, D. (forthcoming b). Representationalism and sensory modalities: an argument for intermodal representationalism. *American Philosophical Quarterly*.
- Bourget, D. and Mendelovici, A. (2014). Tracking representationalism. In Bailey, A., editor, *Philosophy of Mind: The Key Thinkers*, pages 209–235. Continuum.
- Byrne, A. (2001). Intentionalism defended. *Philosophical Review*, 110(2):199–240.
- Chalmers, D. (2015). Panpsychism and panprotopsychism. In Alter, T. and Nagasawa, Y., editors, *Consciousness in the Physical World: Essays on Russellian Monism*. Oxford University Press.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, D. J. (2002). Does conceivability entail possibility? In Gendler, T. S. and Hawthorne, J., editors, *Conceivability and Possibility*, pages 145–200. Oxford University Press.
- Chalmers, D. J. (2003a). Consciousness and its place in nature. In Stich, S. P. and Warfield, T. A., editors, *Blackwell Guide to the Philosophy of Mind*, pages 102–142. Blackwell.
- Chalmers, D. J. (2003b). The content and epistemology of phenomenal belief. In Smith, Q. and Jokic, A., editors, *Consciousness: New Philosophical Perspectives*, pages 220–72. Oxford University Press.
- Chalmers, D. J. (2004). The representational character of experience. In Leiter, B., editor, *The Future for Philosophy*, pages 153–181. Oxford University Press.
- Chalmers, D. J. (2007). Phenomenal concepts and the explanatory gap. In Alter, T. and Walter, S., editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press.
- Chalmers, D. J. (2009). The two-dimensional argument against materialism. In McLaughlin, B. P. and Walter, S., editors, *Oxford Handbook to the Philosophy of Mind*. Oxford University Press.

- Chalmers, D. J. and Jackson, F. (2001). Conceptual analysis and reductive explanation. *Philosophical Review*, 110(3):315–61.
- Churchland, P. S. (1986). *Neurophilosophy: Toward A Unified Science of the Mind-Brain*. MIT Press.
- Crane, T. (2003). The intentional structure of consciousness. In Smith, Q. and Jokić, A., editors, *Consciousness: New Philosophical Perspectives*, pages 33–56. Oxford University Press.
- Crick, F. and Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex*.
- Descartes, R. (1641). *Meditations on First Philosophy*. MIT Press.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.
- Eccles, J. C. (1978). *Mind and Brain*. Paragon House.
- Goff, P. (2009). Why panpsychism doesn't help us explain consciousness. *Dialectica*, 63(3):289–311.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(April):127–136.
- Kripke, S. A. (1980). *Naming and Necessity*. Harvard University Press.
- Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly*, 64(October):354–61.
- Lockwood, M. (1989). *Mind, Brain, and the Quantum*. Oxford University Press.
- Lycan, W. G. (1996). *Consciousness and Experience*. MIT Press.
- Maxwell, G. (1971). Structural realism and the meaning of theoretical terms. *Minnesota Studies in the Philosophy of Science*, 4:181–192.
- Maxwell, G. (1979). Rigid designators and mind-brain identity. *Minnesota Studies in the Philosophy of Science*, 9:9.
- Mendelovici, A. (forthcoming). *The Phenomenal Basis of Intentionality*. Oxford University Press.

- Pautz, A. (2009). A simple view of consciousness. In Koons, R. C. and Bealer, G., editors, *The Waning of Materialism: New Essays*, pages 25–66. Oxford University Press.
- Penrose, R. (1994). *Shadows of the Mind*. Oxford University Press.
- Rey, G. (1995). Toward a projectivist account of conscious experience. In Metzinger, T., editor, *Conscious Experience*, pages 123–42. Ferdinand Schoningh.
- Russell, B. (1927). *The Analysis of Matter*. London: Kegan Paul.
- Schlosshauer, M. (2005). Decoherence, the measurement problem, and interpretations of quantum mechanics. *Rev. Mod. Phys.*, 76:1267–1305.
- Seager, W. E. (1995). Consciousness, information, and panpsychism. *Journal of Consciousness Studies*, 2(3):272–88.
- Seager, W. E. (2002). Panpsychism. In McLaughlin, B., Beckermann, A., and Walter, S., editors, *Stanford Encyclopedia of Philosophy*. Stanford: Metaphysics Research Lab.
- Speaks, J. (2010). Attention and intentionalism. *Philosophical Quarterly*, 60(239):325–342.
- Speaks, J. (2015). *The Phenomenal and the Representational*. Oxford University Press.
- Stapp, H. P. (1996). The hard problem: A quantum approach. *Journal of Consciousness Studies*, 3(3):194–210.
- Stoljar, D. (2001). Two conceptions of the physical. *Philosophy and Phenomenological Research*, 62(2):253–281.
- Strawson, G. (1994). *Mental Reality*. MIT Press.
- Strawson, G. (2003). Real materialism. In Antony, L. M., editor, *Chomsky and His Critics*, pages 49–88. Blackwell Publishing.
- Strawson, G. (2006). Realistic monism - why physicalism entails panpsychism. *Journal of Consciousness Studies*, 13(10-11):3–31.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.