

Is Consciousness Fundamental?

Philip Goff

This paper is forthcoming in Rabin, Gabriel (Ed.) *Grounding and Consciousness*, Oxford University Press.

Abstract: In part I, I argue that the direct access each person has to their own consciousness reveals substantial information about the essential nature of conscious experience, and that this should be taken as a fundamental source of data in metaphysical enquiry. In part II, I argue – partly based on this data – that consciousness is not grounded in purely structural facts, of the kind physical science provides. I then explore alternative options for accommodating consciousness in our theory of reality, ultimately supporting a non-reductive or semi-reductive form of panpsychism.

I – The Reality of Consciousness

The general public in the UK has no idea that something called ‘metaphysics’ goes on in most philosophy departments in the country. When I tell people about this fact, explaining that metaphysics is philosophical enquiry into the nature of reality – rather than something to do with crystals – the standard reaction is bafflement as to how one could hope to find out what reality is like without doing experiments. This bafflement is understandable. How on earth is one supposed to find out about reality just by sitting in your armchair thinking about stuff? What are the *data* that are supposed to inform metaphysical enquiry?

Of course, analytic metaphysics is not completely divorced from science, and metaphysicians regularly appeal to empirical data in their arguments. It is also common to make recourse to a priori intuitions. I’m not here going to raise objections to either of these common practices. Instead, I want to suggest there is another source of data that ought to inform metaphysical enquiry, one that is unfortunately somewhat neglected in contemporary analytic philosophy. I have in mind our introspective judgements concerning our own conscious experience.

It’s plausible that these judgements enjoy a special kind of justification that can warrant a much higher credence than is permissible with respect to our perceptual judgements about our bodies or the world around us. A symptom of this is that skepticism about one’s own experiences – *am I really feeling this pain?* – is much harder to entertain with than skepticism about the external world – *is there really a table there?*

However, this special justification is not much use to metaphysics if introspection reveals to us nothing about the *essential nature* of our conscious experiences. Many physicalists believe that attending, say, to one’s pain, and reflecting of it in terms of how it feels, reveals almost nothing of its essential nature. When I introspect, on this view, I sort of *point* inside myself, but without having a clue what I’m pointing at. It turns out I’m pointing at a brain state, but this fact was hidden from introspection.

I think this view is self-evidently false. I appreciate that’s not much of an argument for those who’ve reflected on the matter carefully and don’t find it self-evidently false. But perhaps a little guided meditation will persuade those who don’t have a dog in a race. I would prompt such guided meditation as follows:

In so far as I'm focused on the *feeling* of pain, the object of my thought is essentially defined by how people feel when they feel that way. This is just true by definition (as Humpty Dumpty wisely observed, I can mean what I want by my words). But what is it for someone to feel a certain way? Surely (the 'surely' operator expresses an invitation to agree with me about what I take to be self-evident) you know the answer to that question when you know how it feels to feel that way. And when I attend to my pain, I do know how it feels. In knowing how it feels, I thereby know what it is to feel that way.

This may sound trivial spelt out like that – I hope so, because trivial propositions are true propositions. But my physicalist opponents typically deny it.¹ They tend to accept the first bit: that pain is essentially defined by how it feels (to this extent, few are willing to depart from the Gospel according to Kripke).² However, strange as it sounds, they think that knowing *what it's like* to feel a given pain tells you nothing about *what it is* to feel that way. They have to say this because, on their view, experiences have a purely physical nature – to have a given experience is to instantiate some specific physical or functional state – and yet our introspective attention to our experiences reveals nothing of that physical or functional nature (if only neuroscience were so easy!).

If you found yourself persuaded by the above meditation (and please do take a little time to think about it before making a judgement...reflecting on your intuitions should take longer than Googling...), you accept something like the thesis that's become known as 'Revelation'.³ The basic idea of Revelation is that introspective attention to our experiences reveals to us the essential nature of those experiences.

The thesis of Revelation needs to be put carefully. Often our introspective judgement are rough and ready – *that bloody headache again* – and I wouldn't want to assert that these kinds of sketchy judgements perfectly capture every nuance of the experiences they refer to. And sometimes we make mistakes about the character of our experiences, e.g. hastily think we're experiencing orange when we're in fact experiencing red.

In a slightly different context, David Chalmers (2002) tries to get around these kinds of worries by focusing on *direct phenomenal concepts*. The phrase 'phenomenal concept' is the technical term for the kind of concept we form when we think about experience in terms of how it feels or what it's like. Whilst many such concepts are based on prior experiences, *direct* phenomenal concepts are those whereby the entire content of the concept is formed through attention to a current experience. In other words, I attend to a particular experience I'm currently having, and I just think *that* – where I'm attending to what it's like to have the experience. Chalmers thinks introspective judgements formed on the basis of direct phenomenal concepts enjoy a privileged epistemic status.

In my first book *Consciousness and Fundamental Reality* I went along with this. But more recently, however, some doubts have set in. Some of these doubts are based on my experience of meditation. When you first start meditating, you feel like you're just attending to the breath, and taking in what's there to be experienced. In Chalmers' terms, it seems like you're sustaining a direct phenomenal concept of your breath experience. But a couple of hours in, it becomes apparent that, in fact, you're projecting onto your breath experience *an idea of the breath*, a sort of caricature of the breathing experience – based on your many previous such experiences – as opposed to just taking in the highly

¹ McLaughlin 2001 and Papineau 2006.

² Kripke 1980

³ For defences of Revelation, see Nida-Rumelin 2007, Goff 2011, 2015b, 2017, Liu 2019, 2020b, 2021. The word 'revelation' was first used by Mark Johnson (1992) to describe Galen Strawson's (1989) view of colour experience, but was later applied to consciousness by Daniel Stoljar (2006).

specific breath experience you're currently having. This raises for me a doubt as to whether we ever in fact have any perfectly direct phenomenal concepts, or whether we're always projecting onto our experience ideas based on past experiences.

Here's another source of doubt. I'm rubbish at tasting. I suspect I'm, as it were, shortsighted taste-wise. I put too much salt on everything. But I hear when statistically normal tasters go on wine-tasting courses, they learn about all sort of hidden depths of their wine tasting experiences, hints of parsley and the like. Suppose a novice attends to their experience of a specific merlot at the start of a wine-tasting course, and endeavours to form a direct phenomenal concept of their gustatory experience. Later, at the end of the course, that same individual attends to an experience of the same merlot and, again, endeavours to form a direct phenomenal concept. Presumably, there will be all sorts of subtle aspects of the experience that they hadn't picked up at the start of the course. Has the experience changed? Maybe. But equally maybe (if that's a phrase, which it's not), the experience is pretty much the same, but the wine-tasting course has enabled this individual to apprehend much more about what it's like to have the experience. If that's right, then the first application of what seemed to be a direct phenomenal concept didn't reveal the complete essence of the experience it denoted, as it left out all the subtle hints of parsley the novice could enjoy only after receiving their diploma.

I propose getting around this by defining Revelation as follows:

Revelation (refined definition): For any phenomenal concept C referring to an experience E, in so far as C accurately represents what it's like to have E, to that extent C reveals the essence of E.

Again, this might sound like rendering trivial the Revelation thesis. And again, if that strikes you as trivial, then I'm happy, as trivial propositions are true. But, as before, this is not trivial, as many (perhaps most) physicalists deny it. A standard physicalist position is that even when a phenomenal concept very accurately captures the character of an experience, it still doesn't tell us much or anything about the essential nature of that experience.

We now have a definition of Revelation and an assertion that its truth is self-evident (if that's not convincing, I've also argued (Goff 2015b, 2017) that Revelation offers the best explanation of the special justification enjoyed by our introspective judgements concerns our experiences). If Revelation is true, then introspection yields a crucial source of data for metaphysics. Our experiences are 'subjective', in the sense that they are experiences. But they are also part of objective reality: it's an objective fact that I'm feeling an ache in my knee right now (it's not, because I'm just making it up for the sake of an example, but you get the idea). And therefore, if Revelation is true, there's a bit of reality that (a) I know for certain exists, (b) I have direct access to the essential nature of. This is metaphysical gold dust! I really believe that future metaphysicians will be baffled that late 20th/early 21st century metaphysicians didn't make more of this datum.⁴

What can be established from this course of data? For one thing, the falsity of physicalism (at least on one understanding of what physicalism is). This is the focus of the next section.

II – The Anti-Physicalist Argument

Physicalists believe that all facts, including the facts about consciousness are grounded in, and thereby necessitated by, the physical facts. As Frank Jackson (2000) puts it: any minimal physical

⁴ Wild assertions about what future historians will think is a rhetorical ploy I make great use of.

duplicate of our world – a world physically just like ours and with nothing else in it – would be a duplicate per se of our world. To put the point in terms of Kripke's (1980) theological metaphor, physicalists believe that in creating all universe, all God had to do was to fix the fundamental physical facts: the arrangements and physical properties of particles and fields and so on (or strings or wave functions, of whatever turn out to be the entities that feature in fundamental physics). These fundamental physical facts, for the physicalist, necessitate all the other facts, including the existence and character of our conscious experience.

An important question in assessing the plausibility of physicalism is whether the putative grounding relation between the physical and the experience is *intelligible*. We can make clear the notion of intelligibility I have in mind by bringing in Laplace's demon, the imaginary super intelligence dreamed up by the 19th century French physicist Pierre-Simon Laplace. Laplace (1825/1951: 4) imagined the demon to have complete knowledge of the fundamental physical facts *at a time*, and an intellect 'vast enough to submit these data to analysis.' On that basis, Laplace believed that the demon would be able to deduce the entire past and future of the universe, 'for such an intellect nothing would be uncertain and the future just like the past could be present before its eyes.'

The probabilistic nature of (some interpretations of) quantum mechanics make it less clear that the demon could perfectly deduce the past and future. So let us give our demon a bit more to work with, allowing them complete knowledge of the fundamental physical facts across all of time. Let us also make explicit that the demon has unlimited reasoning capacities. Finally, we can attribute to Laplace's demon a phenomenal concept of every actually instantiated conscious experience.⁵

Let us stipulate that the question 'Is the grounding relation between the physical and the consciousness facts intelligible?' is synonymous with the question 'Could Laplace's demon (as defined above) work out all of the consciousness facts – including the character of each conscious being's experiences – on the basis of the physics facts?' There are two standard answers physicalists give to this question:

Intelligibility via Structure – Conscious experiences are *structurally defined*, i.e. where *structure* is what can be captured in a purely causal or mathematical vocabulary. By seeing that the physical facts realise the relevant structures, Laplace's demon can deduce everything there is to know about my experiences, including their phenomenal character.⁶

Unintelligible Grounding – The grounding connection between the physical facts and consciousness is inherently unintelligible: even with knowledge of all the physical facts, and a phenomenal concept of each actually instantiated experience, Laplace's demon would not be able to work out who has consciousness and which experiences belong to whom.⁷

Let us take each of these options in turn. The problem with *Intelligibility via Structure* is that it is just not plausible that phenomenal concepts are purely structural concepts. A colourblind scientist who knows everything about red experiences that can be captured with causal and mathematical concepts would not thereby know what it's like to have a red experience. Conscious experiences have a lot of structure, but you cannot capture the content of phenomenal thought – the kinds of

⁵ It is commonly held that to have a phenomenal concept of a given experience one needs to actually have or have had the experience, in which case we can suppose that the demon has had all actually instantiated experiences.

⁶ Armstrong 1968, Dennett 1991, Lewis 1994.

⁷ Loar 1990/97, Papineau 2002, Diaz-Leon 2010

thought we have when we think about our experiences in terms of what it's like to have them – in purely structural terms.

Are there not other ways of making the psycho-physical grounding connection intelligible? The facts of physical science are *structural facts*, which is to say they are expressed in a vocabulary of only mathematical and causal or nomic concepts ('nomic' meaning to do with laws of nature). Now of course it doesn't follow that there is nothing more to *the physical world* than pure structure. Panpsychists, for example, hold that there is more to the physical world than physics reveals to us (more on this presently). But for the moment I want to consider the view that the experiential facts are grounded in the kinds of facts physics reveals to us: the 'physicSal' facts, to use Galen Strawson's (2006) term. Given that the physicSal facts are purely structural, the only way a proposition asserting the grounding of the experiential in the physicSal could be a priori knowable is if the phenomenal concepts involved in that proposition were purely structural.

What about *Unintelligible Grounding*? This might be plausible if Revelation were false. Suppose phenomenal concepts revealed nothing of the essential or accidental nature of their referents; they just kind of 'point.' Now imagine Laplace's demon has a phenomenal concept of a red experience, and is surveying the physical processes in the brain, wondering where to locate that experience. Given that the concept just refers without providing any information about the experience, Laplace's demon will have nothing to go on. If we assume the falsity of Revelation, it makes perfect sense that the truths concerning the grounding of consciousness in the physical would not be a priori knowable.

But now suppose Revelation is true: the demon's phenomenal concept of a red experience completely reveals the essential nature of that experience. And recall the setup that the demon has a complete understanding of the fundamental physical facts across space and time, those facts which the physicalist thinks necessitate the facts about consciousness. If the physicalist continues to accept *Unintelligible Grounding*, then, I suggest, they're accepting a very odd kind of brute necessity. This kind of brute necessity would surprise even an omnipotent creator.⁸ We can imagine our creator bringing into existence all the physical truths, and being surprised to find that they necessitate the existence of consciousness.

It is one thing to hold that contingent truth are brute, in the sense of not being a priori knowable; we couldn't work out a priori the laws of physics, for example. And even when it comes to necessary truths, it makes sense that these truths might not be a priori knowable if the concepts involved in them don't reveal the essences of the properties they denote. For example, the concept of 'water' doesn't reveal that water is essentially H₂O, and hence the necessary truth that 'water is H₂O' is not a priori. All of this is consistent with the following principle:

Grounding Intelligibility – For any proposition N of the form <X is grounded in Y>, where it is a priori (for someone who fully understands the proposition) what it would be for N to be true, N is true if and only if N is a priori knowable.⁹

Why accept Grounding Intelligibility? At this point we get very close to philosophical bedrock. To my mind, the whole project of metaphysics is about trying to provide an intelligible story of reality. You have to take some facts as brute – and these are the facts that determine the 'price' of your theory. And then if you don't want to pay more, you have to put the work in to show that everything else

⁸ The traditional omniscient God knows all truths and hence would, by definition, know that the physical truths necessitate the experiential truths (if that is indeed a truth). Hence why I focus on an omnipotent creator without specifying that they're omniscient.

⁹ Something like this principle is defended in Goff 2017 and Goff 2019c.

you believe in could be intelligibly derived from the facts you are paying for (i.e. the facts that are fundamental on your theory).

For any entity we believe in, we can always ask the question 'Would Laplace's demon have been able to deduce the existence of this thing from other facts (e.g. from the micro-level facts)?', and what answer we give on that ought to make a difference. It's plausible that Laplace's demon would have been able to deduce the existence of a table-shaped object in my kitchen from their knowledge that there are particles arranged table-wise in my kitchen, and that seems to give a clear sense in which the table in my kitchen is not ontologically additional to – does not cost more than – the micro-level facts that obtain in my kitchen. Whereas if the demon couldn't deduce my consciousness from the physical facts of my brain, then my consciousness is an ontological addition relative to the physical facts of my brain and so we have to spend more money if we want to accommodate my consciousness into our metaphysical picture of reality.

Jonathan Schaffer's (2017) has recently defended a view he calls 'physicalism' which conflicts with the letter but not the spirit of Grounding Intelligibility. On this view, all grounding relations are mediated by grounding principles, or 'laws of metaphysics.' We can make this clear by comparing Schaffer's 'physicalism' to David Chalmers' naturalistic dualism. For Chalmers, whilst zombie worlds are perfectly possible, we are fortunate that in the actual world there are contingent psycho-physical laws of nature which ensure that certain brain processes causally produce consciousness. Schaffer's view is exactly the same, except that he thinks the relevant laws are *metaphysically necessary*, rather than contingent, and result in the relevant brain processes *grounding* rather than *causing* consciousness.

I don't see the theoretical advantage of insisting that these laws are necessary and thus making some necessary truths brute, surprising even to an omnipotent being. And to say the productive relationship here is 'grounding' rather than 'causation' seems to me a distinction without a difference. Nonetheless, Schaffer is not really positing an *unintelligible* grounding relation, as the existence of the grounding principles renders the grounding relation intelligible: Laplace's demon could deduce my consciousness from my brain processes in conjunction with the grounding principles. For this reason, Schaffer's view seems to me closer to dualism than physicalism. We don't want to get too hung up on semantics; suffice to say, Schaffer's view is not in the class of views I'm arguing against in this section of the paper.¹⁰

The final question we need to ask is whether the following is true:

Psycho-Physical Revelation – For any proposition N of the form <Q is grounded in P>, where Q is an experiential fact (conceived of under ideal phenomenal concepts) and N is a physical fact, it is a priori (for someone who fully understands the proposition) what it would be for N to be true.

In line with the refined definition of Revelation articulated above, to the extent that the phenomenal concepts involved in <Q> fail to fully and correctly articulate what it's like to have the experience they denote, to that extent <Q> will fail to fully and correctly articulate what it is to have

¹⁰ Part of Schaffer's reason for calling his view 'physicalism' is that he thinks there's nothing special about the grounding of the mental in the physical; for any macro-level entity, its grounding in the physical needs to be facilitated by substantive grounding principles. I disagree with this. I don't think we need fundamental grounding principles to bridge the gap between the fact that people are partying and the fact that there is a party: all it is for there to be a party is for there to be people partying. Similarly, the existence of, say, rocks requires nothing more than that there are particles arranged rock-wise. The mind is a rare case in that its existence requires more than there being particles arranged 'mind-wise.'

those experiences, and to that extent it will not be a priori what it is for <Q> to be true. Given that <Q> is part of the overall proposition <Q is grounded in P>, if it's not fully a priori what it is for <Q> to be true, then it's not fully a priori what it is for <Q is grounded in P> to be true; in other words, Psycho-Physical Revelation will be false. We can get around this problem via the specification that <Q> involves 'ideal' phenomenal concepts, i.e. phenomenal concepts that perfectly articulate the character of the experiences they denote. We can clearly make sense of such ideal phenomenal concepts, even if the real-world phenomenal concepts employed by human beings always fall slightly short of this ideal.

So much for <Q>. What about <P>? I am very sympathetic to the idea that our purely structural physical concepts fail to fully capture the nature of the physical. Indeed, the panpsychist view I advocate below relies on this. However, it seems tautological that our purely structural physical concepts fully capture the nature of the *structural*. Or rather, in so far as our physical concepts are correctly applied, to that extent they perfectly capture the nature of the structure. In other words, so long as our focus is on the 'physicSal (rather than the 'physical' in some more liberal use of that word) the situation with <P> is the same as the situation with <Q>: our physical concepts fail to reveal the essential nature of the physical world only in so far as they are misapplied. We need not worry about potential misapplication of concepts when evaluating Psycho-Physical Revelation, as we are concerned not with propositions that are true but rather with *what it would be* for any such proposition to be true (in particular whether that information is a priori accessible). We can conclude, therefore, that, for any true grounding claim of the form <Q is grounded in P>, it will be a priori what it would be for <P> to be true.

The final bit of the grounding claim we need to consider is the term expressing the grounding relation. Is it a priori what it is for one fact to be grounded in another? One might think we latch onto the grounding relation 'out there' in reality, without grasping its nature. But such a view would seem to be inconsistent with the grounding connection being intelligible, in the sense articulated above. Consider the grounding relation that obtains between the fact that there are people partying at a certain location and the fact that there is a party at that location. Suppose I knew that there was some relation that obtained between these facts, a relation I called 'grounding,' but without knowing anything about its nature. In this case, that grounding relation in question would be something of a 'black box' for me: I know it takes the input of people partying and it outputs a party, but without any understanding of what is involved in these facts being related in this way. Yet the grounding relation between these two facts clearly is intelligible: if I know there are people partying at Sara's house, I thereby know a priori that there is a party at Sara's house. It is self-evident that the grounding fact (the fact that people are partying) *must* stand in the ground relation to the grounded fact (the fact that there is a party), and it's hard to see how this could be self-evident if we lacked a priori understanding of the nature of the grounding relation.¹¹

We have now established that all three components of <Q is grounded in P> (for any grounding claim of this form) transparently reveal what they are claiming about reality. Of course, one cannot assume that a whole shares the properties of its parts. But where a proposition fails to transparently reveal what it would be for it to be true, that is because some component of it defers to reality, as it were, in fixing its satisfaction-conditions, which thereby (at least in part) ensures that reality is deferred to in order to fix the truth-conditions of any proposition in which that concept features. For example, humans fixed the meaning of the concept <water> by pointing to water, with the implicit intention

¹¹ Perhaps we have some but not complete a priori understanding of the grounding relation. I find it hard to see what such a view would say about how we fix reference to the grounding relation, such that we have some but not complete access to its nature. If such a view comes along, I'll consider it then.

that 'water' is *that stuff, whatever it is*. If none of the components of <Q is grounded in P> do this, then there is no way in which the truth-conditions of <Q is grounded in P> could be fixed by anything other than the a priori accessible content of that proposition. I conclude, therefore, that Psycho-Physical Revelation is true.

Bringing all of this together, we can formulate the following argument against the view that experiential facts are grounded in the facts of physical science:

An Argument Against PhysicSalism

Grounding Intelligibility – For any proposition N of the form <X is grounded in Y>, where it is a priori (for someone who fully understands the proposition) what it would be for N to be true, N is true if and only if N is a priori knowable.

Psycho-Physical Revelation – For any proposition N of the form <Q is grounded in P>, where Q is an experiential fact (conceived of under ideal phenomenal concepts) and N is a physicSal fact, it is a priori (for someone for fully understands the proposition) what it would be for N to be true.

Explanatory Gap – For any proposition N of the form <Q is grounded in P>, where Q is an experiential fact and N is a physicSal fact, N is not a priori knowable.

Conclusion – For any proposition N of the <Q is grounded in P>, where Q is an experiential fact and P is a physicSal fact, N is not true.

III – Alternatives to PhysicSalism

The above argument, if it works, shows that experiential facts cannot be grounded in purely structural facts. There are two alternative possibilities:

- Our conscious experiences are fundamental, in the sense of not being grounded in anything.
- Our conscious experiences are grounded in more-than-structural facts, i.e. facts that may involve structure but cannot be exhaustively captured in structural terms.

Dualism is one form of the first possibility. Perhaps the physical and the mental are both fundamental but radically distinct aspects of reality, closely intertwined through their intimate causal relationships with each but with neither metaphysically dependent on the other. Such a view faces much discussed worries pertaining to parsimony and causal closure (Papineau 2000).

The second option has the potential to avoid these concerns. Suppose there is a more-than-structural grounding base which grounds both the experiential facts and the structural facts of physics. If that grounding base consists of only one kind of property/individual, then we avoid the parsimony worries with dualism. And if the grounding the mental in the structural avoids causal closure worries, then grounding both the mental and the structural in some more-than-structural base has the potential to also remove causal closure worries.

What could this more-than-structural grounding base consist in? The two most discussed options are:

Reductive Panpsychism – Certain facts concerning very rudimentary forms of consciousness ground both the structures of physics and the experience of humans and animals.¹²

Noumenalism – We don't know what the more-than-structural grounding base of reality is, and perhaps can never know.¹³

It is well understood how a panpsychist grounding base could ground the structures of physics.¹⁴ So long as the posited rudimentary conscious entities interact in the ways specified by the laws of physics, then they will thereby ground the existence of the mathematico-causal structures specified by physics. We can't ground consciousness in the structures of physics, but we know how to ground the structures of physics in consciousness.

The problem for reductive panpsychism is that there are considerable doubts that facts about the experience of humans and animals could be intelligibly grounded in facts concerning distinct, more simple conscious entities. For any collection of very simple conscious entities, no matter how complex their interactions, it seems that we can conceive of those simple conscious entities existing in the absence of any *further* conscious entities, such as conscious humans.¹⁵ If we can't ensure an intelligible grounding relationship between facts about the consciousness of simple conscious entities and the facts about human/animal consciousness, then reductive panpsychism faces similar challenges to those faced by physicalism.

Noumenalism avoids this problem, but only because we don't have a clue what the alleged grounding base is, and so we're not in a position to judge whether or not we would get an intelligible grounding relation. Our ignorance renders us free to imagine that knowledge of those basic facts would yield a satisfactory explanation of consciousness. Of course, the downside is that this is a very unsatisfying solution. Still, if we are confident enough that the physical world is causally closed, and that neither physicalism nor reductive panpsychism can deliver the goods, we may be forced to this unsatisfying view as the most likely to be true.

My own view is that reflection on what it is to be a conscious subject gives us grounds for thinking that all conscious subjects are fundamental, including human and non-human animal conscious subjects. I have argued for this at greater length elsewhere and will just give the bare bones of the argument here.¹⁶ A macro-level entity can be intelligibly grounded in certain micro-level entities only if the macro-level entity admits of analysis in more fundamental terms. Consider the example of a party. All it is for there to be a party is for there to be people partying. This metaphysical analysis of partyhood allows to make sense of how facts about the existence of parties can be grounded in facts about the existence of partyers. But our concept of a conscious subject does not admit of analysis of the form 'All it is for there to be a conscious subject is for there to be such and such entities with such and such properties.' For there to be a conscious subject is just for there to be a thing that has

¹² For an overview of recent literature on panpsychism see Goff et al 2017/2022. For volumes on the topic, see Brüntrup & Jaskolla 2016; Seager 2019; Goff & Moran 2022.

¹³ There is also the option of *panqualityism* (Coleman 2012, 2014) the view that fundamental properties are *unexperienced qualities*. I argue in Goff 2017: 6.2.3 that panqualityism faces the same challenges as reductive panpsychism in accounting for the grounding of human and non-human animal consciousness.

¹⁴ This is the approach known as 'Russellian monism,' is rooted in Russell's *Analysis of Matter*, and further developed in Feigl 1967, Maxwell 1979, Lockwood 1989, Strawson 1994, 2003, 2006, Chalmers 1996, 2015, Griffin 1998, Stoljar 2001, Pereboom 2011, 2015 and Goff 2015a, 2017, 2019a, Mørch 2024. Alter & Nagasawa 2015 is a volume of essays on the view.

¹⁵ This argument is developed in more detail in Goff 2009, 2017, Chalmers 2016.

¹⁶ Goff 2015a, 2016, 2019b, 2020a.

subjective experience; nothing less, nothing more. What it is for there to be a conscious subject cannot be expressed in more basic terms.¹⁷

What about causal closure worries? I'm not convinced that there is an empirical reason to accept causal closure; particularly, as Hedda Hassel Mørch has pressed (2014: 3.1, 2023: 4.4.3), if we're focusing on microphysical causal closure, as opposed to a more general thesis of physical causal closure. The experiments on which we base fundamental physics typically involve very small numbers of particles. It is not practically possible to test whether the Born rule of quantum mechanics could perfectly predict the objective probabilities of particle locations in a complex living system like a brain. And I don't think we know enough about the workings of the brain of any conscious organism to know whether all of its behaviour is reducible to known chemistry.¹⁸ The physicist Sean Carroll (2021) has argued that it would be challenging to reconcile currently known physics with forms of fundamental causation that only arise in complex systems.¹⁹ That may be so. But the reality of macro-level conscious subjects is a hard datum for metaphysics. And to the extent that that datum leads us to fundamental macro-level causation – to make sense of mental causation – then the task of metaphysics is to accommodate that fundamental macro-level causation. It may be hard, but who said nature would make it easy for us? Furthermore, Jonathan Emery (2025) has recently responded to Carroll's argument, articulating a place for fundamental macro-level mental causation compatible with current physics.

Is this dualism? Not necessarily. A non-reductionist panpsychist holds that macro-level conscious subjects – such as the conscious minds of humans and non-human animals – and causally brought into and sustained in existence by forms of consciousness at the level of fundamental physics. This view may not have any advantages over dualism in terms of causal closure, as both involve fundamental macro-level causation (although see Mørch 2014, Ch. 3, 2024: 4.4.3 for arguments that non-reductionist panpsychism fares better with causal closure than dualism). However, it is superior to dualism in terms of parsimony, as reductive panpsychism postulates only one kind of fundamental individual and property: conscious subjects and their experiences. Why postulate two kinds of fundamental thing when you only need one?

There is also the possibility of a semi-reductionist form of panpsychism (Goff 2024, Goff and Roelofs forthcoming). Perhaps my conscious mind is fundamental, but it shares some or all of its experiences with the micro-level entities which sustain its existence. One day more detailed understanding of the brain and the neural correlates of biological consciousness may make some such semi-reductionism empirically plausible.

Overall, then, non-reductionist, or semi-reductionist, panpsychism seems to me most likely to be true, on the grounds that it is simplest coherent hypothesis.

References

Alter, T. & N. Yujin (Eds.) (2015) *Consciousness in the Physical World*, Oxford University Press.

Armstrong, D. (1968) *A Materialist Theory of Mind*, Cambridge: Cambridge University Press.

¹⁷ For a contrary view defending reductive panpsychism, see Roelofs 2019.

¹⁸ Matthew Cobb's excellent book *The Idea of the Brain* most convinced me of this (although he certainly isn't a supporter of my philosophical views!).

¹⁹ I've also had many conversations about this with Sean Carroll, on my Mind Chat podcast, and some of this is expressed in Sean's 2021 paper, which I respond to in Goff 2021. These papers are printed in the volume Goff and Moran 2022.

- Brüntrup, G. & Jaskolla, L. (2016) *Panpsychism*, Oxford University Press.
- Carroll, S. M. (2021) 'Consciousness and the Laws of Physics,' *Journal of Consciousness Studies* 28: 9-10, 16-31.
- Chalmers, D. J. (1996) *Consciousness and its Place in Nature*, Oxford University Press.
- Chalmers, D. J. (2002) 'The content and epistemology of phenomenal belief,' in A. Jokic & Q. Smith (Eds.) *Consciousness: New Philosophical Perspectives*, Oxford University Press, 220—72.
- Chalmers, D. J. (2015) 'Panpsychism and protopsychism,' in Alter & Nagasawa 2015.
- Chalmers, D. J. (2016) 'The combination problem for panpsychism,' in Brüntrup, G. & Jaskolla 2016.
- Coleman, Sam (2012) 'Mental chemistry: Combination for panpsychists,' *Dialectica* 66: 1, 137-66.
- Coleman, Sam (2014) 'The real combination problem: Panpsychists, micro-subjects and emergence,' *Erkenntnis* 79: 1, 19-44.
- Dennett, D. (1991) *Consciousness Explained*, Boston: Little Brown.
- Diaz-Leon, Esa (2010) 'Can phenomenal concepts explain the epistemic gap?' *Mind* 119: 476, 933-51.
- Emery, Jonathan (2025) *The Empirical Evidence Against Interactionist Dualism*, PhD thesis, King's College London.
- Feigl, H. (1967) 'The "mental" and the "physical": The essay and a postscript,' University of Minnesota Press.
- Goff, Philip (2011) 'A posteriori physicalists get our phenomenal concepts wrong,' *Australasian Journal of Philosophy*, 89: 2, 191-209.
- Goff, P. (2015a) 'Against constitutive forms of Russellian monism,' in T. Alter & Y. Nagasawa (Eds.) *Russellian Monism*, Oxford University Press.
- Goff, P. (2015b) 'Real acquaintance and physicalism,' in P. Coates & S. Coleman (Eds.) *Phenomenal Qualities: Sense, Perception, and Consciousness*, Oxford University Press.
- Goff, P. (2017) *Consciousness and Fundamental Reality*, Oxford University Press.
- Goff, P. (2019a) *Galileo's Error: Foundations for a New Science of Consciousness*, Pantheon.
- Goff, P. (2019b) 'Grounding, analysis and Russellian monism,' in S. Coleman (Ed.) *The Knowledge Argument Then and Now*, Cambridge University Press.
- Goff, P. (2019c) 'Essentialist modal rationalism,' *Synthese* 198 (Suppl 8): 2019-2027.
- Goff, P. (2020a) 'Cosmopsychism, micropsychism and the grounding relation,' in W. Seager (Ed.) *The Routledge Handbook of Panpsychism*.
- Goff, P. (2020b) 'Revelation, Consciousness+ and the Phenomenal Powers View,' *Topoi* 39: 1089–1092.
- Goff, P. (2021) 'Putting Consciousness First: Replies to Critics,' *Journal of Consciousness Studies* 28: 9-10, 289-328.

Goff, P. (2024) *How Exactly Does Panpsychism Help Explain Consciousness?* *Journal of Consciousness Studies* 31: 3, 56-82.

Goff, P. & Moran, A. (Eds.) (2022) *Is Consciousness Everywhere? Essays on Panpsychism*, Academic Imprint.

Goff, P. & Roelofs, L. (forthcoming) 'In Defence of Phenomenal Sharing,' in L. Bugnon & M. Nida-Rumelin (Eds.) *The Phenomenology and Self-Awareness of Conscious Subjects*, Routledge.

Griffin, D. R. (1998) *Unsnarling the World-Knot: Consciousness, Freedom, and the Mind-Body Problem*, University of California Press.

Jackson, F. (2000) *From Metaphysics to Ethics*, Oxford University Press.

Johnston, M. (1992) 'How to speak of the colours,' *Philosophical Studies*, 68: 221-63.

Kripke, S. (1980) *Naming and Necessity*, Harvard University Press.

Laplace, P. (1825/1951) *A Philosophical Essay on Probabilities*, translated into English from the original French 6th ed. by Truscott, F.W. and Emory, F.L., Dover Publications.

Lewis, D. (1994) 'Reduction of mind,' in S. Guttenplan (Ed.) *Companion to the Philosophy of Mind*, Oxford: Blackwell.

Loar, B. (1990/97) 'Phenomenal states,' originally published in J. Tomberlin (Ed.) *Philosophical Perspectives 4: Action Theory and Philosophy of Mind*, Atascadero, CA: Ridgeview; reprinted in substantially revised form in N. Block, O. Flanagan, & Guuven Guuzeldere (Eds.) 1997, *The Nature of Consciousness: Philosophical Debates* 1: 113-29.

Lockwood, M. (1989) *Mind, Brain and Quantum*, Oxford University Press.

Maxwell, G. (1979) 'Rigid designators and mind-brain identity,' in C. W. Savage (Ed.) *Perception and Cognition*, University of Minnesota Press.

McLaughlin, B. P. (2001) 'In defence of new-wave materialism: A response to Horgan and Tienson,' in C. Gillett & B. Loewer (Eds.) *Physicalism and its Discontents*, Cambridge University Press.

Mørch, H. H. (2014). *Panpsychism and Causation: A New Argument and a Solution to the Combination Problem*, Ph.D. thesis, University of Oslo.

Mørch, H. H. (2023) *Non-Physicalist Theories of Consciousness*, Cambridge: Cambridge University Press.

Liu, M. (2019) 'Phenomenal Experience and the Thesis of Revelation,' in D. Shottenkirk, M. Curado & S. S. Gouveia (Eds.) *Perception, Cognition and Aesthetics*, Routledge, 227-251.

Liu, M. (2021) 'Revelation and the intuition of dualism,' 199, 11491–11515.

Liu, M. (2020) 'Explaining the intuition of revelation,' *Journal of Consciousness Studies*, 27: 5-6, 99-107(9).

Nida-Rumelin, M. (2007) 'Grasping phenomenal properties,' in Alter, T. & Walter, S. (Eds.) *Phenomenal Concepts and Phenomenal Knowledge*, Oxford University Press.

Papineau, D. (2000) 'The rise of physicalism,' in C. Gillett & B. M. Loewer (Eds.), *Physicalism and its Discontents*, Cambridge University Press.

- Papineau, D. (2002) *Thinking about Consciousness*, Oxford: Oxford University Press.
- Papineau, D. (2006) 'Comments on Galen's Strawson: Realist Monism: Why physicalism entails panpsychism,' *Journal of Consciousness Studies*, 13: 10-11. 100-9.
- Pereboom, Derk (2011) *Consciousness and the Prospects for Physicalism*, Oxford University Press.
- Pereboom, Derk (2015) 'Consciousness, physicalism, and absolutely intrinsic properties,' in Alter & Nagasawa 2015.
- Roelofs, L. (2019). *Combining Minds: How to Think about Composite Subjectivity*, New York: Oxford University Press.
- Schaffer, Jonathan (2017) 'The Ground Between the Gaps,' *Philosophers' Imprint* 17.
- Seager, W. E. (2020) (Ed.) *Routledge Handbook of Panpsychism*, Routledge.
- Stoljar, D. (2006) *Ignorance and Imagination*, Oxford University Press.
- Strawson, G. (1989) 'Red and "red,"' *Synthese* 78: 193-232.
- Strawson, G. (1994) *Mental Reality*, MIT Press.
- Strawson, G. (2003) 'Real materialism,' in L. M. Antony (Ed.) *Chomsky and His Critics*, Blackwell, 49-88.
- Strawson, G. (2006) 'Realistic monism - why physicalism entails panpsychism,' *Journal of Consciousness Studies* 13: 10-11, 3-31.
- Goff, Philip, Seager, William, & Allen-Hermanson, Sean (2017/2022) 'Panpsychism,' *Stanford Encyclopedia of Philosophy*.