

# Chapter 7

## Orthodox Property Dualism + The Linguistic Theory of Vagueness = Panpsychism

Philip Goff

By ‘consciousness’ I mean *the property of being a thing such that there’s something that it’s like to be that thing*. The meaning of this rather cumbersome phrase can be illustrated with reference to our commonsense beliefs about what things have the property it denotes. According to common sense, there’s something that it’s like for a rabbit to be cold, or to be kicked, or to have a knife stuck in it. In contrast, there’s nothing that it’s like for a table to be cold, or to be kicked, or to have a knife stuck in it. There’s nothing that it’s like *from the inside*, as it were, to be a table (according to common sense). Consciousness, as I will understand it, is the property of *having an inner life* of some kind or other; a property ordinary opinion supposes to be confined to the biological realm.

There are a number of powerful arguments in the literature – I will focus on the zombie-conceivability argument and the knowledge argument – which have the conclusion that consciousness is a non-physical feature of reality.<sup>1</sup> Call these arguments ‘the standard arguments’. A sizeable minority of philosophers (i) accept the soundness of the standard arguments, and so take consciousness to be a non-physical feature of reality, (ii) nonetheless take consciousness to be a property of physical objects rather than immaterial substances, a basic property which arises from physical properties in accordance with fundamental psycho-physical laws of nature. Call such philosophers ‘orthodox property dualists’. The purpose of this

---

<sup>1</sup>Chalmers (1996, 2002) and Jackson (1982).

P. Goff (✉)

Department of Philosophy, University of Liverpool, 7 Abercromby Square,  
Liverpool L69 7WY, UK

e-mail: [philgoff1@googlemail.com](mailto:philgoff1@googlemail.com)

paper is to argue that orthodox property dualism, in conjunction with the linguistic theory of vagueness, implies *panpsychism*: the view that consciousness is ubiquitous in nature.<sup>2</sup>

Of course, one might accept this conclusion and go a number of ways with it. Depending on the strength of a property dualist's antecedent commitments, accepting my argument might lead her to embrace panpsychism, or to embrace metaphysical vagueness, or to look hard again for a flaw in the standard arguments. I will not be exploring any of these options in what follows. Nonetheless, I take it to be philosophically significant in itself that the conjunction of two popular views has such a surprising implication.

The argument will proceed in three stages. In Sect. 7.1, I will argue that the orthodox property dualist is committed to two theses concerning the *concept* of consciousness: *conceptual dualism* and *phenomenal transparency*. In Sect. 7.2, I will argue that the orthodox property dualist who accepts the linguistic theory of vagueness, because of her commitment to *phenomenal transparency* and *conceptual dualism*, must accept *phenomenal precision*: the thesis that it can never be vague whether or not a given thing is conscious. In Sect. 7.3, I argue from *phenomenal precision* to *panpsychism*. In Sect. 7.4, I will support my case with some methodological remarks.

## 7.1 Conceptual Dualism and Phenomenal Transparency

### 7.1.1 Conceptual Dualism

Each of the standard arguments kicks off with an epistemic premise: zombies are conceivable, Mary learns something new when she leaves her room. For each of these epistemic premises, accepting its truth commits one to the following principle:

*Conceptual Dualism*: The physical facts do not entail the phenomenal facts, i.e. there is no way of moving a priori from knowing the kind of things physics has to tell us about the world to knowing what conscious states there are, or indeed whether there are any conscious states.

If the physical facts entailed the phenomenal facts, then zombies would be inconceivable, and pre-liberated Mary would know what it was like to see red. The orthodox property dualist, by definition, accepts the soundness of the standard arguments, and therefore is committed to *conceptual dualism*.

---

<sup>2</sup>Note that my argument is not primarily aimed at forms of anti-physicalism other than property dualism, such as Russellian monism (although I think the argument has some force against Russellian monism as I explain in footnote 14), nor against property dualists who do not take the standard arguments to be sound.

### 7.1.2 Phenomenal Transparency

Each of the standard arguments begins with an epistemic premise. Each of the standard arguments tries to derive from its epistemic premise a metaphysical conclusion: zombies are possible, the physical description of reality is incomplete. Doing this requires a commitment to the following principle:

*Phenomenal transparency:* The concept consciousness (I will refer to concepts with underlined words) reveals the nature of consciousness, i.e. it is a priori (for someone possessing the concept consciousness, in virtue of possessing that concept) what it is for something to be conscious.

Spelling out this principle, and why the orthodox property dualist is committed to it, requires a bit of work.

It is plausible that concepts denoting properties come divided up into two categories: transparent and opaque. A transparent property concept reveals the nature of the property it denotes, in the sense that it is a priori (for someone who possesses the concept, in virtue of possessing the concept) what it is for an object to have that property.<sup>3</sup> To put it another way, a transparent property concept reveals what is ascribed in an application of the concept. An opaque property concept reveals nothing of what it is for an object to instantiate its referent<sup>4</sup> (I develop this framework for thinking about concepts in more detail in Goff (2011, MS) and Goff and Papineau (forthcoming)).

The best way to clarify and make the case for this distinction is by giving examples. Suppose David's favourite property is Euclidean sphericity, but I am blissfully unaware of this joyous fact. Now consider two ways in which I might think about Euclidean sphericity. I might think of it as *David's favourite property*, where I use that description as a rigid designator. Alternately I might think of it in geometrical terms, as the property of *being a thing with all points on its surface equidistant from its centre*. There is a clear sense in which, when I think of Euclidean sphericity as *David's favourite property*, I don't understand its nature. I have no idea what it is for something to instantiate 'David's favourite property', or as we might simply put it *I have no idea what David's favourite property is*. In contrast, when I think about the same property in geometrical terms I do understand its nature. I know what it is for an object to be spherical: it's for it to have all points on its surface equidistant from its centre. The concept Euclidean sphericity is transparent; the concept David's favourite property (rigidly designated) is opaque.

<sup>3</sup>Concept *C* renders fact *F* a priori if it is metaphysically possible for someone to know *F* in virtue of possessing *C*, without relying on any empirical information beyond what is required to possess *C*.

<sup>4</sup>An opaque concept may (but may not) reveal accidental properties of the property it denotes, e.g. it is plausible to think that the concept *being water* reveals that in the actual world the property denoted realises the property of being the colourless, odourless stuff in oceans and lakes. I call an opaque concept which reveals accidental properties of the referent which uniquely identify it in the actual world 'mildly opaque' (see footnote 5).

If consciousness is taken to be opaque, there is no way of moving beyond the epistemic premise of any of the standard arguments. Consider the zombie conceivability argument. For those physicalists – probably currently the majority – who accept that zombies are conceivable, the challenge is to explain why the conceivability of zombies does not entail their genuine possibility. If it is an option to hold that consciousness is opaque, it is obvious what the physicalist can say:

The concept consciousness denotes a purely physical or functional property – that’s why zombies are impossible – but because consciousness is opaque, it’s not a priori that consciousness denotes a purely physical or functional property – that’s why zombies are conceivable.

If we allow that consciousness is opaque, the conceivability of zombies has no metaphysical significance.<sup>5</sup>

Consider the knowledge argument. The challenge for the physicalist is to say what Mary learns upon liberation. If it is possible that the concept consciousness, and consequently our concepts of its determinates, i.e. specific modes of consciousness, are opaque, it is obvious what the physicalist can say:

---

<sup>5</sup>On Chalmers 2D semantic framework (1996, 2002, 2009) the primary intention of each term/concept is a priori evaluable (without this, the move in his two-dimensional argument from the conceivability of a state of affairs to its genuine possibility at some world considered as actual would be implausible). He also holds, as most people do, that the primary and secondary intentions of consciousness are the same (which justifies the move in the two-dimensional argument from the possibility of zombies at some world considered as actual to the possibility of zombies at some world considered as counterfactual). It follows that consciousness has an a priori evaluable secondary intention, which is equivalent to its being transparent, see Nida-Rümelin 2007 for a detailed analysis of this (where I talk about concepts ‘revealing the nature’ of properties, she talks of concepts ‘affording a grasp’ of properties). Thus, Chalmers’ two-dimensional argument against materialism, at least in its standard form, is dependent on the thesis that consciousness is transparent.

Chalmers does claim that the two-dimensional argument goes through without the premise that the primary and secondary intentions of consciousness are identical, as he believes that the conceivable truth of  $\langle P \& \sim Q \rangle$  – where P is the complete physical truth about our world and Q is some arbitrary phenomenal truth about our world – entails that the primary intention of that proposition is true at some world W, and given that W is a minimal physical duplicate of our world but not a duplicate simpliciter, physicalism must be false. The idea is that Q might be what I have called elsewhere ‘mildly opaque’, i.e. does not reveal the nature of its referent, but reveals accidental features of the referent which uniquely identify it in the actual world. In this case, although W is physically indiscernible from us, it lacks certain properties, i.e. the properties which uniquely identify Q in the actual world, and hence physicalism is false. Similarly to the case of translucency discussed in *Aside: Why not translucency?* (below), if the orthodox property dualist wanted to say that consciousness is mildly opaque, I would focus on the accidental features of consciousness that consciousness does reveal to us the nature of, rather than focusing on consciousness itself, and argue for the conclusion that those accidental features are ubiquitous in nature. However, I don’t know of any anti-physicalists who do take this approach; it would mean distinguishing the properties we use to think about consciousness, i.e. the property of being a thing such that there’s something that it’s like to be that thing, from the essential nature of consciousness itself.

Liberated Mary gains a phenomenal concept which denotes a purely physical or functional property, but is conceptually novel for her because it is opaque, and hence it is not a priori that it denotes a purely physical or functional property. Therefore, when she leaves the room, Mary does not become acquainted with a new feature of reality, but rather finds a new way of thinking about a feature she already knew about in her room.

In the case of each of the standard arguments, a move beyond the merely epistemic is premised on denying that consciousness is opaque. The orthodox property dualist, i.e. the property dualist who is a property dualist on account of the standard arguments, is committed to *phenomenal transparency*.<sup>6</sup>

### 7.1.3 *Aside: Why Not Translucency?*

I have divided up property concepts into the transparent and the opaque. But these categories do not seem to be exhaustive. Why think that each concept reveals either *all* or *nothing* of the nature of its referent? There seems room for the category of *translucent concept*, where a property concept is translucent if it reveals *some but not all* of the nature of the property it denotes, i.e. something but not everything of what it is for an object to have that property is a priori knowable (for someone possessing the concept, in virtue of possessing the concept). Is it open to the orthodox property dualist to take consciousness to be translucent: revealing some but not all of the nature of consciousness?

I take it that if a property concept is translucent, then the property it denotes is complex, involving within itself a number of aspects. At least one aspect will be denoted transparently, i.e. its nature will be a priori accessible, and at least one aspect will be denoted opaquely, i.e. it will be denoted, but its nature not a priori accessible. Take for example the concept being a sphere roughly the same size as the Earth. This concept reveals the nature of one aspect of the property it denotes, i.e. *being a sphere*, but does not reveal the nature of another aspect of the property it denotes, i.e. *being roughly the same size as the Earth*; empirical work must be done to discover the nature of the latter but not the former aspect.

We can thus consider a translucent concept as a composite of two ‘sub-concepts’, one transparent and one opaque. I call the transparent sub-concept the ‘window’ of the whole concept, and the opaque sub-concept the ‘screen’ of the whole concept. In the above example, the concept of being a sphere is the window of the whole concept, whilst the concept being the same mass as the Earth is its screen.

---

<sup>6</sup>Without phenomenal transparency, it is impossible to move beyond the epistemic premise of the standard arguments. But with phenomenal transparency, and the epistemic premise of either of the standard arguments, the falsity of physicalism follows pretty quickly. If consciousness reveals the nature of consciousness, and consciousness does not reveal consciousness to have a physical nature (which follows from the truth of either of the epistemic premises), then consciousness does not have a physical nature.

If someone wants to defend the claim that the concept having an inner life/being something such that there's something that it's like to be it is translucent, then they are obliged to give us an account of how that concept divides into window and screen. Which aspects of the property of having an inner life are a priori accessible, and which do we refer to but not understand without empirical investigation?

As far as I am aware only physicalists have even given such an account of our phenomenal concepts, i.e. the concepts we form when we think about conscious states in terms of what it's like to have them. Robert Schroer, for example, claims that we can know a priori certain facts about the internal structure of conscious states, but not the intrinsic nature of the basic elements in that structure (Schroer 2010). This allows Schroer to combine *conceptual dualism* with an account of phenomenal concepts according to which they reveal significant information about the states they denote. For Schroer, physical states do not entail phenomenal states, as although we know a priori the internal structure of phenomenal states, we don't know a priori whether the elements composing that structure are physical or non-physical (and hence don't know a priori whether the entire state which results is physical or non-physical).

On Schroer's account, although the whole concept is a priori distinct from the physical facts, the window is not: if we knew all the physical facts, we could see that the internal structures connoted by phenomenal concepts are realised in the brain. But for the standard arguments to have force, the window as well as the whole concept must be a priori distinct from the physical facts. Otherwise the physicalist can simply give the following explanation of why zombies are conceivable but not possible (as Schroer in fact does):

For each phenomenal concept, both window and screen denote purely physical or functional properties – that's why zombies are impossible – but because the screen is opaque, it's not a priori that the screen denotes a purely physical or functional property, and hence not a priori that the whole concept denotes a physical or functional property – that's why zombies are conceivable.

For the standard arguments to succeed, there must be at least one aspect of conscious experience which is understood a priori, and which is not entailed by the physical facts. It is difficult to see how an orthodox property dualist might divide the concept of having an inner life into an aspect we transparently understand and an aspect we don't, and indeed difficult to see what their motivation for doing so would be. But if they did divide up the concept into window and screen, we could simply substitute the word 'consciousness' in what follows for 'consciousness\*', defined as 'that aspect of consciousness we understand the nature of a priori'. I will continue to assume that consciousness is transparent, but we can note that if consciousness turns out to be translucent rather than transparent, then my argument is to be read as aiming to show that consciousness\*, rather than consciousness, is ubiquitous in nature.<sup>7</sup>

---

<sup>7</sup>The situation is similar to the case of the imagined – as far as I am aware non-existent – anti-physicalist who wants to claim that we pick out consciousness in virtue of its accidental features, which I discuss in footnote 5.

## 7.2 Orthodox Property Dualism + Linguistic Theory of Vagueness = Phenomenal Precision

Despite well known contemporary defences of epistemic and metaphysical accounts of vagueness, the ‘linguistic theory of vagueness’, i.e. the broad spectrum of views which locate the source of vagueness in language rather than the world, remains probably the most popular approach to dealing with vagueness.

According to the linguistic theory of vagueness, vagueness is the result of *semantic indecision*: for any vague predicate there are multiple ‘sharpenings’ of the predicate, such that the meaning of the predicate does not settle on any of these sharpenings. Consider the vague predicate ‘is tall’. We could stipulate, somewhat arbitrarily, that anything that is exactly 6 ft or taller counts as ‘tall’, and anything shorter is not tall. This is one ‘sharpening’ of the predicate ‘is tall’, that is, one way of making the predicate precise. Alternately, we could stipulate that anything that is exactly 6 ft and 1 in. or taller counts as tall, and anything shorter is not tall. This is an alternative sharpening of ‘is tall’, that is, an alternative possible way of making the predicate precise. The predicate ‘is tall’ is thus associated with a *spectrum of sharpenings*: a range of possible ways of making the predicate precise.<sup>8</sup>

The linguistic theory of vagueness tells us that a vague predicate is vague because no one has bothered to single out one of its sharpenings as the unique meaning of the predicate. To put it metaphorically, the predicate hasn’t made up its mind which of those precise meanings it wants to plump for. Suppose John is a borderline case of tallness. According to the linguistic theory of vagueness, it’s not that in reality there is some fuzzy, indeterminate state of affairs of John’s neither having nor lacking a certain quality. In the world there’s just John with some utterly precise height. It’s *the predicate* that is indeterminate such that there’s no fact of the matter as to whether or not it applies to things with John’s exact height. The indeterminacy is in language rather than the world.

The linguistic theory of vagueness explains the vagueness of a predicate in a way that involves the associated spectrum of sharpenings. Clearly if this kind of explanation is to work, then each vague predicate must be associated with a spectrum of sharpenings. However, the predicate ‘is conscious’ – and hence the concept it expresses – does not seem to be associated with a spectrum of sharpenings, at least not a priori.

This can be a difficult point to get across, because the word ‘consciousness’ is used in lots of different ways by different philosophers and scientists. Sometimes the predicate ‘is conscious’ is used to mean *is aware/cognitively sophisticated to a certain level*, perhaps roughly the level we would be inclined to

---

<sup>8</sup>With some vague predicates, as with ‘is tall’, the sharpenings are determinates of a single determinable. In the case of other vague predicates, e.g. ‘is a religion’, there is a weighted cluster of properties, *involves belief in a supernatural being, involves ritual, involves a moral code*, such that each sharpening involves some of those properties but it is not the case that each sharpening involves all of those properties.

call ‘self-consciously aware’. This does seem to be a notion of consciousness which is associated with a spectrum of sharpenings, ranging from more to less cognitively sophisticated. But that’s not the notion of consciousness we are concerned with. I am using the predicate ‘is conscious’ to mean *has an inner life of some kind or other*, and this doesn’t seem to be a notion of consciousness which is a priori associated with a spectrum of sharpenings. You either have an inner life or you don’t. Of course you can have a richer or a less rich inner life, a more sophisticated or a less sophisticated inner life. But the property of *having an inner life* itself does not present itself to us as one that admits of degree: you either have it or you don’t.

The physicalist wanting to embrace phenomenal vagueness, at least if she is prepared to deny *phenomenal transparency*, need not worry that the sharpenings of consciousness are not available a priori. She can claim that the semantic workings of the concept, and therefore its spectrum of sharpenings, are determined ‘outside the head’. David Papineau, for example, an explicit rejecter of *phenomenal transparency*,<sup>9</sup> denies that the semantic workings of consciousness – constituted of causal or teleological facts – are a priori accessible. Those semantic workings, according to Papineau, leave it indeterminate whether the concept picks out the capacity for higher-order judgement, or the physical basis for that capacity in humans.<sup>10</sup> There is thus a recognisable sense in which consciousness – and hence the predicate ‘is conscious’ – has (at least) two sharpenings: (A) the capacity for higher-order judgement, (B) the physical basis of higher-order judgement in humans. On sharpening (A) silicon duplicates of humans count as ‘conscious’, on sharpening (B) they don’t. Papineau does not take (A) and (B) to be a priori accessible: the semantic workings of consciousness are not a priori accessible, and so neither are the more referentially precise versions of those semantic workings.<sup>11</sup>

Nothing I have said above casts doubt on Papineau’s view, or anything like it. But notice that it assumes the falsity of *phenomenal transparency*, at least if we are assuming the truth of the linguistic theory of vagueness. According to the linguistic theory of vagueness, what is ascribed in the application of a given vague predicate is to be understood in terms of the predicate’s indeterminacy over its sharpenings.

---

<sup>9</sup>In his 2006 Papineau gives an explicit denial of *phenomenal transparency*.

<sup>10</sup>Papineau 2002, ch. 7. In fact, Papineau is open to the possibility of conscious states which cannot be thought about, and because of this ends up thinking that the concept consciousness is indeterminate such that on one sharpening it refers to attention, on one sharpening it refers to pre-attention, and on one sharpening it refers to the property of being material! It is an under-emphasised implication of this (I have confirmed with Papineau in conversation that he embraces this implication), that there is no fact of the matter as to whether or not panpsychism is true, just as there is no fact of the matter as to whether I am tall. On one sharpening of consciousness, the table and the pillar of salt are conscious, on another sharpening they are not. It is ironic that Papineau’s denial of transparency, which allows him to escape the argument for panpsychism given in this paper, gets him in the end to panpsychism (at least on one legitimate sharpening of consciousness).

<sup>11</sup>It is because of this option, open to the a posteriori physicalist like Papineau, of claiming that the semantic workings of consciousness are outside of what is a priori accessible to the concept user, that I reject the kind of argument Michael Antony (2006) gives for the non-vagueness of consciousness.

Assuming the truth of this view, if what is ascribed in the application of a given vague predicate is a priori knowable, then the sharpenings of that predicate must be a priori knowable. In other words, if the linguistic theory of vagueness is true, then the sharpenings of a transparent predicate must be a priori knowable. If the orthodox property dualist wants to claim that consciousness is associated with a spectrum of sharpenings, whilst remaining faithful to the linguistic theory of vagueness, then, given her commitment to *phenomenal transparency*, she is obliged to hold that these sharpenings are a priori accessible.

There seems to me only one even vaguely plausible proposal as to what the spectrum of sharpenings a priori associated with consciousness is: that which would be offered by the analytic functionalist (even this proposal does not seem very plausible to me, but then that is because I don't find analytic functionalism very plausible). If the predicate 'is conscious' is a functional or behavioural predicate, then presumably it is associated a priori with a spectrum of sharpenings, which can be captured with a fine grained enough functional description.<sup>12</sup> But of course the orthodox property dualist, given her commitment to *conceptual dualism*, is obliged to deny that the predicate 'is conscious' is a functional predicate. The functional and behavioural states of an organism are entailed by the physical facts about that organism. Therefore, if 'is conscious' were a functional or behavioural predicate, then it too would be entailed by the physical facts, contrary to *conceptual dualism*. Putting the analytic functionalist's proposal on one side, there just doesn't seem to be another candidate for being the spectrum of sharpenings a priori associated with consciousness.

Perhaps it might be objected that the sense of consciousness can be sharpened, but that we lack the necessary concepts to grasp the resulting sharper concept. By analogy, it might seem initially plausible that someone might possess the concept phenomenal red, without possessing a concept of any more specific phenomenal shade of red. Such a person would possess a concept, the sense of which can be sharpened, and yet be unable to sharpen it. Perhaps this is the situation we are in with respect to the general concept consciousness; the sense of the concept can be sharpened, but we lack the concepts required to do it.

However, for the reasons discussed above, assuming the truth of the linguistic theory of vagueness, there couldn't be a transparent vague concept which does not allow a priori knowledge of its sharpenings. According to the linguistic theory of vagueness, what is ascribed in the application of a given vague predicate is to be understood in terms of the predicate's indeterminacy over its sharpenings. Given this, for what is ascribed to be a priori knowable, it must be a priori knowable what the relevant sharpenings are.

There is a sense in which our inability to find sharpenings of consciousness is not *entirely* conclusive evidence that there are no a priori knowable sharpenings

---

<sup>12</sup>On Lewis's kind of materialism (see Lewis 1994) there would be a priori associated with consciousness a spectrum of sharpenings of the property of *being consciousness*, but not of consciousness itself (Lewis takes mental concepts to be flaccid designators).

of consciousness. Certain facts which are rendered a priori knowable by concept *C* may be out of the cognitive reach of a given individual possessing *C*, due to that individual's cognitive limitations.<sup>13</sup> Perhaps one might suppose that if we were better reasoners we would be able to see how to sharpen consciousness. This seems to me an implausible leap of faith. We are not dealing with some difficult mathematics which is beyond our cognitive capacities, but which greater beings than ourselves could deal with. We are dealing with the basic semantic structure of a single concept. If our best efforts to find sharpenings of consciousness do not yield them, then we must suppose that there are no such things, at least not accessible a priori.

I conclude, therefore, because of their commitment to *phenomenal transparency* and *conceptual dualism*, the orthodox property dualist is unable to make sense of consciousness having sharpenings. If she wants to remain faithful to the linguistic theory of vagueness, the orthodox property dualist must hold that consciousness is not vague.

### 7.3 From *Phenomenal Precision* to Panpsychism

In the bible we hear that God turned Lot's wife into a pillar of salt. You get the impression that it happened pretty quickly, but let's suppose that in fact God did it in really small stages: He took Lot's wife, made a slight adjustment to one fundamental particle, a slight adjustment to another fundamental particle, and so on until He had a pillar of pure salt.

Had God gone about it this way, the result would be a temporally continuous series, with Lot's wife at one end, a pillar of salt at the other, and in between a series of objects such that any two objects next to each other in time differ at most by a slight adjustment of a fundamental particle.

Here's a common sense assumption:

*Commonsense Assumption:* Lot's wife is conscious and a pillar of salt is not conscious.

It follows from *Commonsense Assumption* that we have consciousness at one end of the series but not the other; somewhere along the series consciousness disappears. If it could be vague whether or not a given thing is conscious, then presumably there would be borderline cases along the series, where it is vague whether or not we have case of consciousness. But assuming *phenomenal precision*, the cut off point must be utterly sharp. Somewhere along the line there must be two objects, next to each other in time, differing only by a slight adjustment to a fundamental particle, such that one but not the other is conscious. This leads us to the following implausible consequence:

---

<sup>13</sup>See footnote 3 for my definition of a priori knowability.

*Implausible Consequence*: The fundamental psycho-physical laws which specify the physical conditions nomologically sufficient for consciousness are utterly precise, in the sense that the slightest adjustment to the smallest particle can make the difference between whether or not a macroscopic object is conscious.

Why is *Implausible Consequence* implausible? Consider the following analogy. Imagine one day I am blowing up a blue balloon. I blow it up about two thirds of the way when suddenly, to my surprise, the balloon turns pink! In shock, I let the balloon deflate. I try blowing it up again, and find that, at exactly the same point, the balloon turns pink. I experiment with a number of balloons from the same packet but find that the effect is not repeated. Much experimenting later, I discover that the following is a fact about our universe:

*Random Fact*: When a blue balloon is (i) made from three specific kinds of elastic, A, B and C, such that there is 42 % of A, 38 % of B, and 20 % of C, (ii) has a certain thickness, precise to 1,000,000,000th of a millimetre, (iii) is blown up such that it's diameter has a certain length, precise to 1,000,000,000th of a millimetre, the balloon turns pink.

The hypothesis that *Random Fact* constitutes a basic law of nature is extraordinary. Were we to discover that *Random Fact* obtains in our world, we would be extremely reluctant to take it as a fundamental law, and would try to find a way of explaining its obtaining in terms of more general laws, ones which did not involve such arbitrarily precise values. Of course it is not *inconceivable* that such a law obtains: there is an extremely strange possible world governed by such a law. The hypothesis that such a law obtains is not necessarily false, but is extremely theoretically implausible. It is rational to avoid such a hypothesis if at all possible.

But if the supposition that *Random Fact* constitutes a basic law of nature is to be avoided, then so much, much more so is *Implausible Consequence*. A law *L* specifying that physical conditions *P* are sufficient for macroscopic consciousness, where *P* are utterly precise down to slightest change in the smallest particle, would involve such arbitrarily precise specifications – many times more so than those involved in *Random Fact* – that it would be crazy to suppose that *L* was brute. *Implausible Consequence* is to be avoided at all costs.<sup>14</sup>

---

<sup>14</sup>This argument is aimed at orthodox property dualists, whom I have stipulated to hold that consciousness is a fundamental feature of reality arising in accordance with basic psycho-physical laws of nature. But not all anti-physicalists take consciousness to be a fundamental feature of reality. Many Russellian monists (Russell 1927; Feigl 1958/1967; Maxwell 1979; Lockwood 1989; Chalmers 1996; Griffin 1998; Stoljar 2001) take phenomenal properties be realised in *proto-phenomenal* properties, certain qualities of physical objects which are not themselves phenomenal properties, but are somehow intrinsically suited to constitute phenomenal properties (clearly, our grasp of such qualities is frustratingly meagre). Perhaps the Russellian monist could hold that the conditions sufficient for consciousness are utterly precise, but that this fact is explained in terms of some more fundamental laws involving protophenomenal properties, laws which do not involve such arbitrarily precise specifications. However, even on the supposition that consciousness is not fundamental, it is still pretty implausible to suppose that a slight adjustment to a single fundamental particle – one of countless billions – in the brain could make the difference between the whole brain having or lacking the determinable property of consciousness. So I am inclined to think that

But, as I hope to have shown, *Implausible Consequence* follows from the conjunction of *Commonsense Assumption* and *phenomenal precision*: *Commonsense Assumption* entails that somewhere along the series consciousness disappears, *phenomenal precision* entails that the disappearance of consciousness must be sharp rather than vague. If we want to reject *Implausible Consequence*, then we must reject (at least) either *Commonsense Assumption* or *phenomenal precision*. Given that the orthodox property dualist is committed to *phenomenal precision*, she must reject *Commonsense Assumption*: she must hold either *that neither Lot's wife nor the pillar of salt are conscious* or *that both Lot's wife and the pillar of salt are conscious*. Given her realism about consciousness, the orthodox property dualist is hardly going to go for the former disjunct. Therefore, she is obliged to think that both Lot's wife and the pillar of salt are conscious.

Of course it's not going to end there. We could take any pair of macroscopic objects such that common opinion takes the former but not the latter to be conscious, and do the same thing. To return to the example used at the start of the paper, we might imagine God turning a rabbit into a table, particle by particle, and by a similar chain of reasoning get to the conclusion that the table is conscious. We quickly end up with panpsychism: the view that consciousness is ubiquitous throughout nature.<sup>15</sup>

In setting up the thought experiment I have implicitly assumed unrestricted composition, such that none of the changes God makes to the particles which initially compose Lot's wife results in those particles ceasing to compose anything. Let's entertain the supposition that composition is restricted, such that when the particles are arranged Lot's wife-wise they compose, but when they are arranged pillar of salt-wise they fail to compose; somewhere along the series we cease to have a composite object.

Call the time when the particles stop composing 'C', and the time when the particles stop *phenomenally composing*, that is composing a conscious object, 'P'. Let us consider in turn the supposition (i) that C precedes P (ii) that C is simultaneous with P, (iii) that P precedes C.

Supposition (i) is impossible. It cannot be the case that C precedes P, for any time at which the particles phenomenally compose is a time at which the particles compose.

Supposition (ii) implies that C is a precise time, given that P is a precise time (assuming *phenomenal precision*). The supposition that C is a precise time entails a sharp cut off point between the particles composing and the particles ceasing

---

the considerations outlined here have some force against the Russellian monist, even though the argument is primarily aimed at, and has more force against, the orthodox property dualist.

<sup>15</sup>Throughout the thought experiment I have, for simplicity, assumed that there are fundamental particles, and have spoken of time as though it were ultimately composed of indivisible moments. Neither of these simplifications is essential to the argument. We might instead suppose that God makes a slight adjustment to a sub-atomic particle every 100,000,000,000,000th of a second. Even if there are no fundamental particles, it is still implausible that the fundamental psycho-physical laws of nature are precise such that a slight adjustment of a sub-atomic particle can make the difference between the presence and absence of macroscopic consciousness.

to compose: a slight adjustment of a fundamental particle makes the difference between the particles composing and failing to compose. But this leads to:

*Implausible Consequence\**: The mereological laws which specify the physical conditions sufficient for particles to compose are utterly precise, in the sense that the slightest adjustment to the smallest particle can make the difference between the presence and absence of macroscopic composition.

However, *Implausible Consequence\** is just as implausible as *Implausible Consequence*.<sup>16</sup> It is just as implausible to suppose that there are basic mereological laws involving such arbitrarily precise specifications, as it is to suppose that there are basic psycho-physical laws involving such arbitrarily precise specifications. It is just as implausible to suppose that there are sharp cut off points between macroscopic composition and its absence as it is to suppose that there are sharp cut off points between macroscopic phenomenal composition and its absence. We must reject the supposition that *C* is a precise time, and hence the supposition that *C* is simultaneous with *P*.

Finally, let us consider supposition (iii). Given the implausibility of sharp cut off points between macroscopic composition and its absence, we must suppose that *C* is a vague time. So we are supposing that at some precise time *P* the particles stop phenomenally composing, and at some later vague time *C* the particles stop composing altogether. I don't think this is a plausible supposition, as I hope to demonstrate in what follows.

At the first moment after *P*, call it '*P + 1*', there must be a definite fact of the matter as to whether the particles phenomenally compose (assuming *phenomenal precision*).<sup>17</sup> Given that the particles definitely phenomenally compose at *P*, it is implausible to suppose that they definitely do not phenomenally compose at *P + 1* – this would lead to *Implausible Consequence* – therefore at *P + 1* the particles must definitely phenomenally compose. And if the particles definitely phenomenally compose at *P + 1*, then they definitely compose at *P + 1*.

But now consider the second moment after *P*, call it '*P + 2*'. There must be a definite fact of the matter at *P + 2* whether or not the particles phenomenally compose. Given that they phenomenally composed at *P + 1*, they must phenomenally compose, and hence compose, at *P + 2*, on pain of the truth of *Implausible Consequence*. We could keep doing this for every subsequent moment until we get to the particles arranged pillar of salt-wise, which entails that there is no moment along the series at which the particles stop composing, i.e. *C* does not exist. Supposition (iii) cannot be sustained once we have signed up to *phenomenal precision*.

Thus, once we have committed to *phenomenal precision*, we cannot plausibly hold that any of the adjustments God makes result in the particles failing to compose. We now have a complete argument, not only for panpsychism, but also for unrestricted phenomenal composition, and hence for unrestricted composition,

<sup>16</sup>A similar claim is argued in Sider (2001), 120–134, a strong influence on this argument.

<sup>17</sup>Those who take time to be infinitely divisible may substitute 'moment' for '100,000,000,000,000th of a second', see footnote 15.

at least regarding macroscopic objects. All combinations of particles numerous enough to be arranged macroscopic-wise phenomenally compose, and hence all such combinations of particles compose.

Why do I make the qualification that phenomenal composition is unrestricted ‘regarding macroscopic objects’? I have been implicitly supposing in the above thought experiments that the number of particles remains unchanged in these imagined transformations of woman to pillar of salt, or rabbit to (presumably quite small) table. But what if God took a conscious being and annihilated one particle a time, until only one particle remained? Is the orthodox property dualist obliged to think a single particle has an inner life?

It seems to me that the argument still has force when we are dealing with objects composed of very high numbers of particles. For a conscious object composed of seven billion particles, it is implausible to suppose that the psycho-physical laws are precise such that the removal of a single one of those seven billion particles could render it non-conscious. But it is not clear to me that the argument has force when we are dealing with objects composed of small numbers of particles. The smaller the number of particles required for consciousness, the less implausibly arbitrary the values involved in the psycho-physical laws, e.g. it is not implausible to suppose that the basic psycho-physical laws specify that at least four particles are required for phenomenal composition.<sup>18</sup>

The orthodox property dualist, then, is not obliged to subscribe to *unrestricted phenomenal composition*, but only to *unrestricted phenomenal composition at the macroscopic level*. We thus end up with a very different kind of panpsychism to that defended by contemporary panpsychists such as Galen Strawson<sup>19</sup> and Sam Coleman.<sup>20</sup> These panpsychists warrant the name in virtue of holding that *the fundamental constituents of reality* are conscious, but are reluctant to attribute consciousness to inanimate macroscopic objects. Given the vagueness of the boundary between the animate and the inanimate, and given the commitment to *phenomenal precision* that I would argue the commitment to the soundness of the standard arguments forces upon these panpsychists, the considerations I have outlined above put severe strain on this kind of view.<sup>21</sup>

---

<sup>18</sup>For a similar reason I believe Sider’s ‘vagueness argument’ for unrestricted composition is inconclusive. It does not seem implausible to me to suppose that the basic laws of mereology specify that at least four particles are required for composition. Sider’s argument gives us strong reason to think that *macroscopic* composition is unrestricted, but has no force when applied to cases at the fundamental level involving a small number of particles.

<sup>19</sup>Strawson (2006).

<sup>20</sup>Coleman (2006, 2009).

<sup>21</sup>The argument of this section is heavily influenced by the Lewis/Sider ‘vagueness argument’ for unrestricted composition, see Lewis (1986, 221–213) and Sider (2001, 120–134).

## 7.4 Common Sense and Serious Metaphysics

I would like to finish by strengthening my case with some methodological considerations. One might think that the case I have made is less than conclusive, as the orthodox property dualist can always avoid panpsychism without giving up on the linguistic theory of vagueness by going for *Implausible Consequence*. *Implausible Consequence* is in itself a very unattractive option, but, when the alternative is conscious pillars of salt, one might be forgiven for suddenly finding it attractive.

Even if this thought is right, we still have an interesting result. We have the orthodox property dualist facing a difficult choice between deeply implausible fundamental laws, metaphysical/epistemic accounts of vagueness, and conscious pillars of salt. But I do want to go further, and to do what I said I would do, which is to argue that orthodox property dualists who are committed to the linguistic theory of vagueness should be panpsychists. In order to do this, I must lessen the theoretical concern regarding panpsychism, which is what I will try to do in what follows.

What is the worry about panpsychism? I don't think it can be a worry about economy. For sure the panpsychist believes in more consciousness than does the average man. But this is at worst a sin of *quantitative* rather than *qualitative* profligacy – postulating more of a kind we already believe in rather than postulating new kinds – and it is generally agreed by metaphysicians that quantitative profligacy is not an especially heinous sin. It is postulating new *kinds* of thing beyond necessity that we need to avoid.

I think the worry with panpsychism is simply that it is so at odds with ordinary opinion. But when you take a step back, it's difficult to see why this consideration should concern the metaphysician. If we're trying to find out the nature of reality as it is in and of itself, why should we care what the average Joe thinks about things? Scientists often tell us weird things about the world. How often do other scientists say, 'Now hold on, Steve, this is getting quite out of kilter with what the average person thinks . . . maybe we should have second thoughts . . .'. Not often. And if fit with common opinion is not a serious consideration in science, it is difficult to see why it should be a serious consideration in metaphysics.

One contemporary metaphysician to have offered an argument for a concern for common sense is David Lewis.<sup>22</sup> I assume that something like Lewis's justification is implicitly guiding the practices of contemporary commonsense-ophile metaphysicians:

. . . it is pointless to build a theory, however nicely systematised it might be, that it would be unreasonable to believe. And a theory cannot earn credence just by its unity and economy. What credence it cannot earn, it must inherit. It is far beyond our power to weave a brand new fabric of adequate theory *ex nihilo*, so we must perforce conserve the one we've got [i.e. the theory that is implicit in common sense] . . . It's not that the folk know in their blood

---

<sup>22</sup>Having 'a concern' for common sense does not render it sacrosanct. Arguably Lewis ends up straying quite far from what would be acceptable to the average Joe.

what the highfalutin' philosophers may forget. And it's not that common sense speaks with the voice of some infallible faculty of 'intuition'. It's just that theoretical conservatism is the only sensible policy for theorists of limited powers...<sup>23</sup>

How do we choose between theories in the sciences? One thing we do is weigh theoretical virtues: where there are empirically equivalent theories, we choose between them on the basis of simplicity, elegance, etc. But of course our primary concern, our starting point for enquiry, is fit with the empirical data. We first turn to the empirical data, and then when we've got everything we can there, we turn to theoretical virtues (no doubt an oversimplification, but it'll do).

How do we decide between theories in metaphysics? Again, one thing we do is weigh theoretical virtues. But, as Lewis says, that can't be the starting point for our enquiry; we can't weave a theory out of elegance, simplicity, etc. We could end up anywhere! So the interesting question is: What constitutes the starting point of metaphysical enquiry? What plays the role in metaphysics that empirical data plays in science?

Lewis, because he doesn't think there's anything better, opts for common sense. The Lewisian method is to start with the theory that is implicit in common sense, and then move beyond that on the basis of theoretical virtues. Crucially, Lewisian metaphysics is built on common sense only because *there isn't anything better*.

But the orthodox property dualist does have something better. The orthodox property dualist claims to have a concept which: (i) transparently reveals the nature of its referent, and (ii) is satisfied. Indeed, I take it that most orthodox property dualists believe that we know with Cartesian certainty that the concept of consciousness is satisfied; each person knows for certain that s/he is conscious. A transparent concept which we know for certain is satisfied amounts to a window onto a bit of the world as it is in and of itself. Much better than common sense!

Unlike Lewis, the orthodox property dualist has no need for common sense; she is able to build metaphysics on much firmer foundations. She might, like Descartes, try to start and finish with *that which cannot be doubted*. However, history is testimony to the failure of Descartes' research project. The orthodox property dualist metaphysician would be better advised to steer a middle way between Descartes and David Lewis. She should follow Descartes in starting with the undoubtable, but follow Lewis in moving beyond the starting point of enquiry by appeal to theoretical virtues. Here's the slogan: Start with *the undoubtable*, then move to *that which the undoubtable renders most probable*<sup>24</sup> (I develop this 'post-Galilean' approach to metaphysics in much more detail in my (MS)).

The only reason a metaphysician need care about common sense is from want of anything better upon which to build metaphysics. But the orthodox property

---

<sup>23</sup>Lewis (1986, 134).

<sup>24</sup>Strictly speaking we have certainty only *that the concept of consciousness is satisfied*. We are not infallible concerning what it takes for the concept to be satisfied (although I take it that we can have strongly justified knowledge about the latter).

dualist has something better: a priori access to the complete nature of a certain feature of reality, i.e. consciousness. The orthodox property dualist should forget about common sense, and embrace conscious pillars of salt.<sup>25</sup>

## References

- Antony, Michael V. 2006. Vagueness and the metaphysics of consciousness. *Philosophical Studies* 128(3):515–538.
- Chalmers, D. 1996. *The conscious mind*. Oxford: Oxford University Press.
- Chalmers, D. 2002. Consciousness and its place in nature. In *Philosophy of mind: Classical and contemporary readings*, ed. D. Chalmers, 247–272. Oxford/New York: Oxford University Press.
- Chalmers, D. 2009. The two-dimensional argument against physicalism. In *Oxford handbook to the philosophy of mind*, ed. Brian P. McLaughlin, 313–339. Oxford: Oxford University Press.
- Coleman, S. 2006. Being realistic: Why panpsychism may entail panexperientialism. *Journal of Consciousness Studies* 13(10–11): 40–52.
- Coleman, S. 2009. The mind that abides. In *Mind under matter*, 83–107. Amsterdam/Philadelphia: John Benjamins Publishing Company.
- Feigl, H. 1958/1967. The ‘mental’ and the ‘physical’. *Minnesota Studies in the Philosophy of Science* 2: 370–497. Reprinted (with a postscript) as *The ‘mental’ and the ‘physical’*. Minneapolis: University of Minnesota Press.
- Goff, P. 2011. A posteriori physicalists get our phenomenal concepts wrong. *Australasian Journal of Philosophy* 89(2): 191–209.
- Goff, P. MS. *Taking consciousness seriously*.
- Goff, P., and D. Papineau. forthcoming. What’s wrong with strong necessities? *Philosophical Studies*
- Griffin, D.R. 1998. *Unsnarling the world-knot: Consciousness, freedom, and the mind-body problem*. Berkeley: University of California Press.
- Jackson, F. 1982. Epiphenomenal qualia. *The Philosophical Quarterly* 32: 127–136.
- Lewis, D. 1986. *On the plurality of worlds*. Oxford: Basil Blackwell.
- Lewis, D. 1994. Reduction of mind. In *Companion to the philosophy of mind*, ed. S. Guttenplan. Oxford: Blackwell.
- Lockwood, M. 1989. *Mind, brain, and the quantum*. Oxford: Oxford University Press.
- Maxwell, G. 1979. Rigid designators and mind-brain identity. *Minnesota Studies in the Philosophy of Science* 9: 365–403.
- Nida-Rümelin, M. 2007. Grasping phenomenal properties. In *Phenomenal concepts and phenomenal knowledge. New essays on consciousness and physicalism*, ed. T. Alter and S. Walter, 307–349. Oxford: Oxford University Press.
- Papineau, D. 2002. *Thinking about consciousness*. Oxford: Oxford University Press.
- Russell, B. 1927. *The analysis of matter*. Repr., New York: Dover, 1954
- Schroer, R. 2010. Where’s the beef? Phenomenal concepts as both demonstrative and substantial. *Australasian Journal of Philosophy* 88(3): 505–522.
- Sider, T. 2001. *Four dimensionalism*. Oxford/New York: Oxford University Press.
- Stoljar, D. 2001. Two conceptions of the physical. *Philosophy and Phenomenological Research* 62: 253–281.
- Strawson, G. 2006. Realistic monism: Why physicalism entails panpsychism. *Journal of Consciousness Studies: Controversies in Science and the Humanities* 13(10–11): 3–31.

---

<sup>25</sup>I would like to thank everyone who contributed at the ‘Mental as Fundamental’ conference in Vienna, and the participants in the second Online Consciousness Conference, especially my two respondents Jonathan Simon and William Robinson.