

Property Dualists Should be Panpsychists

Abstract

This paper attempts to show that property dualists are obliged to hold that it can never be vague whether or not a given thing is conscious. Together with the commonsense assumption that some but not all macroscopic objects are conscious, this entails that the psycho-physical laws governing the emergence of consciousness are utterly precise such that a slight adjustment to a fundamental particle can make the difference between a macroscopic object having or lacking consciousness. I argue that this is a deeply implausible consequence, and hence that the property dualist must reject the common sense assumption that some but not all macroscopic objects are conscious. Given her realism about consciousness, the property dualist is obliged to hold that all macroscopic objects are conscious.

Word count: 8387

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.¹ 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 ‘property dualists’. The purpose of this paper, as the title suggests, is to argue that property dualists so defined are obliged to be *panpsychists*, that is, to hold – in opposition to ordinary opinion – that consciousness is ubiquitous in nature.

Of course, one might accept my conclusion, and go one of two ways with it. The committed property dualist who accepts my conclusion will become a panpsychist, or at least think she ought to. The committed physicalist who accepts my conclusion will no doubt think of this as just one more absurd implication of property dualism, one more reason to reject the view. I will not in this paper be pushing in either of these directions. I merely want to argue for the conditional: if you are a property dualist, then you ought to be a panpsychist.

The argument will proceed in three stages. In section I, I will argue that the property dualist is committed to two theses concerning the *concept* of consciousness: *conceptual dualism* and *phenomenal transparency*. In section II, I will argue that the property dualist, 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

¹ Chalmers 1996, 2002, Jackson 1982.

thing is conscious. In section III, I argue from *phenomenal precision* to *panpsychism*. In Section IV, I will support my case with some methodological remarks.

Section I: Conceptual Dualism and Phenomenal Transparency

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 property dualist, by definition, accepts the soundness of the standard arguments, and therefore is committed to *conceptual dualism*.

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 property dualist is committed to it, requires a bit of work.

It is plausible to think that property concepts come divided 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.² To put it another way, a transparent property concept reveals what is ascribed in an application of the concept. An opaque property concept is one that it is not opaque, that does not reveal the nature of the property it denotes.³

² Concept *C* renders fact *F* a priori iff 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*.

³ 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 in lakes.

The best way to 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 geometry as *David's favourite property*, I don't understand its nature; I have no idea what it is for something to be David's favourite property; I have no idea what David's favourite property is! Whilst 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.

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.⁴

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:

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

⁴ 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 is dependent on the thesis that consciousness is transparent.

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 property dualist is committed to *phenomenal transparency*.⁵

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 iff 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 property dualist to take consciousness to be translucent: revealing some but not all of the nature of consciousness?

I can only make sense of the notion of translucency in the following way. If a property concept is translucent, then the property it denotes is complex, involving within itself a number of aspects. The concept reveals to its possessor the nature of some of those aspects, but not others. A plausible example of this is the concept being a sphere roughly

⁵ 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.

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.

It is difficult to see how we might make sense of the concept having an inner life/being something such that there's something that it's like to be it along these lines. The concept seems too simple; it does not seem to denote multiple aspects such that it would be plausible to say that we understand the nature of some but not others. Perhaps this is not a knock down argument, but in the absence of some proposal as to how to make sense of the consciousness as translucent, I will assume in what follows that it is either transparent or opaque. On the assumption that consciousness is either transparent or opaque, the property dualist is obliged to hold that it is transparent.

Section II: From Property Dualism to Phenomenal Precision

Reflection on cases of vagueness makes it plausible that vagueness is a *semantic* rather than a *metaphysical* phenomenon; that is, gives us reason to sign up to some version of what Theodore Sider has dubbed 'the linguistic theory of vagueness'.⁶ 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.

⁶ Sider 2001: 123-6. Perhaps I am here expressing, without argument, my own intuitions about the most plausible way of accounting for vagueness. However, these intuitions will not play a key role in my argument. Below I consider linguistic, epistemic and metaphysical accounts of phenomenal vagueness.

Consider the vague predicate ‘is tall’, We could stipulate, somewhat arbitrarily, that anything that is exactly 6 feet or taller counts as ‘tall’, and anything shorter is not tall. This is one ‘sharpening’ of the predicate ‘is tall’; one way of making the predicate precise. Alternately, we could stipulate that anything that is exactly six feet and one inch or taller counts as tall, and anything shorter is not tall. This is an alternative sharpening of ‘is tall’; 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.⁷

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.

⁷ 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, e.g. *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.

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 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*,⁸ 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.⁹ 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.

Nothing I have said above casts down 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. Assuming the

⁸ In his 2006 Papineau gives an explicit denial of *phenomenal transparency*.

⁹ 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).

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, the sharpenings of a transparent predicate must be a priori knowable. If the 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.¹⁰ But of course the 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.

¹⁰ 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).

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.

I think reflection reveals the case of prima facie conceivability described above to be incoherent. Our concept phenomenal red does not refer to a specific phenomenal quality. Rather it refers to a number of phenomenal qualities, which all bear a certain resemblance to each other. Therefore, someone could not possess the phenomenal concept phenomenal red without possessing the concept of one of the specific phenomenal qualities to which it refers. Of course one can imagine a person who can introspect a single phenomenal quality which they call 'phenomenal red'. But the concept that person expresses with the phrase 'phenomenal red', call it phenomenal red*, cannot be the same as our concept phenomenal red, for phenomenal red* refers to a single phenomenal quality but phenomenal red does not.

Furthermore, for the reasons discussed above, assuming the truth of the linguistic theory of vagueness, it is difficult to see how there could be a transparent 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, how could what is ascribed be a priori knowable, if it's not 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 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.¹¹ 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 property dualist is unable to make sense of consciousness having sharpenings. If she wants to remain faithful to the linguistic theory of vagueness, the property dualist must hold that consciousness is not vague.

¹¹ See footnote 2 for my definition of a priori knowability.

Perhaps the property dualist might try the following:

In most cases vagueness is a semantic phenomenon, but in the, possibly unique, case of consciousness we have a case of real *metaphysical vagueness*. The world in and of itself comes divided up into things that are conscious, things that aren't conscious, and things in some fuzzy state of being neither conscious nor not conscious.

Careful a priori reflection on consciousness yields a strong modal intuition that this is not possible. Consciousness seems to be an essentially all or nothing, on/off property. It doesn't seem to be the kind of thing whose nature allows for the possibility of fuzzy states of affairs where it's neither had nor not had. The denier of *phenomenal transparency* can happily reject this intuition. If the concept of consciousness is opaque, then any gut feelings we have concerning the essential nature of consciousness are merely reflections of our psychology, not reflections of the real essential nature of consciousness.

But the proponent of *phenomenal transparency* cannot ignore this intuition. For the phenomenal transparentist, *consciousness as we conceive of it* just is *consciousness as it is in and of itself*. If consciousness as we conceive of it is an essentially all or nothing property, it seems to follow that consciousness in and of itself is an essentially all or nothing property (given phenomenal transparency).

Consider the following analogy. Suppose we accept that the ontological argument is sound, such that the concept God must, as a matter of logical necessity, be satisfied. Given this, we cannot very well say, ‘Well God as we conceive of him is all powerful, but God in reality is pretty feeble’. If God represents its object as all powerful, and it’s logically necessary that God is satisfied, it follows that God is in reality all powerful. Similarly, if our concept of consciousness represents consciousness as an essentially on/off property, and that concept is satisfied, it follows that consciousness is an essentially on/off property.¹²

At any rate, if the property dualist wants to believe in metaphysical phenomenal vagueness, then she ought to in some way explain away our modal intuition that such a thing is not possible, and to do so in some way that is consistent with *phenomenal transparency*. I see no way of doing this.

Finally, we can note that the proponent of *phenomenal transparency* is unable to adopt an epistemic view of phenomenal vagueness, for a reason similar to that why she is unable to adopt the Papineau strategy discussed above. If there are aspects of the meaning of consciousness beyond our epistemic ken, it follows that part of what is ascribed in an application of consciousness is beyond our ken, contrary to *phenomenal transparency*.

¹² I take it that the property dualist believes consciousness to be satisfied simpliciter, rather than imperfectly satisfied. For the same reason that it is difficult to make sense of consciousness being translucent, it is difficult to make sense of consciousness being imperfectly satisfied. There don’t seem to be numerous aspects to the concept of having an inner life, such that some but not others might be instantiated.

I conclude, therefore, that the property dualist cannot make sense of phenomenal vagueness, whether of the semantic, epistemic, or metaphysical variety. The property dualist is obliged to be a proponent of *phenomenal precision*.

Section III: 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:

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 third 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 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, and has a certain thickness,

precise to 1,000,000,000th of a millimetre, and is blown up such that its 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.¹³

¹³ This argument is aimed at 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 1926, 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

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 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 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

phenomenal properties (clearly, our grasp of such qualities is frustratingly meagre). Perhaps the Russellian monist could hold that *Implausible Consequence* obtains in our world, but 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 still seems 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 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 property dualist.

the conclusion that the table is conscious. We quickly end up with panpsychism: the view that consciousness is ubiquitous throughout nature.¹⁴

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

¹⁴ 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 are 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. It is still very 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.

point between the particles composing and the particles ceasing 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*.¹⁵ 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

¹⁵ A similar claim is argued in Sider 2002: 120-134, a strong influence on this argument.

don't think we can make sense of this 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*).¹⁶ 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

¹⁶ Those who take time to be infinitely divisible may substitute 'moment' for '100,000,000,000,000th of a second', see footnote 14.

have a complete argument, not only for panpsychism, but also for unrestricted phenomenal composition, and hence for unrestricted composition, 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 small) table. But what if God took a conscious being and annihilated one particle a time, until only one particle remained? Is the 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 specify that at least four particles are required for phenomenal composition.¹⁷

¹⁷ For a similar reason I believe Sider’s ‘vagueness argument’ for unrestricted composition to be inconclusive. It does not seem implausible to me to suppose that the basic laws of mereology specify that at

The 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 Strawson¹⁸ and Sam Coleman.¹⁹ 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.²⁰

Section IV: 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 property dualist can always avoid panpsychism 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.

lest 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.

¹⁸ Strawson 2006.

¹⁹ Coleman 2006, 2009.

²⁰ The argument of this section is heavily influenced by the Lewis/Sider 'vagueness argument' for unrestricted composition, see Lewis 1986: 221-13, Sider 2002: 120-34.

Even if this thought is right, we still have an interesting result. We have the property dualist facing a difficult dilemma between deeply implausible fundamental laws, and conscious pillars of salt. But I do want to go further, to do what I said I would do, which is to argue that property dualists are obliged to be panpsychists. In order to do this, I must lessen the theoretical concern regarding panpsychism. This 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.

This is an under-explored question. One of the few contemporary metaphysicians to have offered an argument for a concern for common sense is David Lewis.²¹ I assume that something like Lewis's justification is implicitly guiding the practices of contemporary commonsense-o-philes:

‘...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 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...’²²

Here's how I interpret Lewis. Compare to science. 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

²¹ 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.

²² Lewis 1986: 134.

data. We first turn to the empirical data, and then when we've got everything we can there, we turn to theoretical virtues (not doubt a horrible 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 property dualist does have something better. The 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 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 property dualist has no need for common sense; she is able to build metaphysics on certain 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 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.*

The only reason a metaphysician need care about common sense is from want of anything better upon which to build metaphysics. But the property dualist has something better: intuitive certainties concerning the nature of consciousness. The property dualist should forget about common sense, and embrace conscious pillars of salt.

References

- Chalmers, D. 1996 *The Conscious Mind*, Oxford: Oxford University Press.
- Chalmers, D. 2002. 'Consciousness and its place in nature', in Chalmers, D. (ed.) *Philosophy of Mind: Classical and Contemporary Readings*, Oxford, New York: Oxford University Press, 247-72.
- Chalmers, D. 2009 'The two-dimensional argument against physicalism', in McLaughlin (ed.) *Oxford Handbook to the Philosophy of Mind*, Oxford: Oxford University Press, 313-39.
- Coleman, S. 2006. 'Being realistic: why panpsychism may entail panexperientialism', *Journal of Consciousness Studies*, 13: 10-11, 40-52.
- Coleman, S. 2009. 'Mind under matter', *The Mind that Abides*, Amsterdam/Philadelphia, John Benjamins Publishing Company, 83-107.

- 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.
- 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', *Philosophical Quarterly*, 32: 127-36.
- Jackson, F. 1986. 'What Mary Didn't Know', *Journal of Philosophy*, 83: 291-295.
- Lewis, D. 1986 *On the Plurality of Worlds*, Oxford: Basil Blackwell.
- Lewis, D. 1994. Reduction of mind. In S. Guttenplan (ed.) *Companion to the Philosophy of Mind*, 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.
- Papineau, D. 2002. *Thinking about Consciousness*, Oxford: Oxford University Press.
- Sider, T. 2001 *Four Dimensionalism*, Oxford, New York: Oxford University Press.
- Stoljar, D. 2001. Two conceptions of the physical. *Philosophy and Phenomenological Research* 62:253-81.
- Strawson, G. 2006. *Realistic monism: why physicalism entails panpsychism*, 13: 10-11, 3-31.