

# The Combination Problem for Panpsychism\*

David J. Chalmers

## 1 Introduction

Panpsychism, the view that fundamental physical entities have conscious experiences, is an exciting and promising view for addressing the mind–body problem. I have argued in “Panpsychism and Panprotopsychism” that it promises to share the advantages of both materialism and dualism and the disadvantages of neither. In particular, it can respect both the epistemological intuitions that motivate dualism and the causal intuitions that motivate physicalism.

Nevertheless, panpsychism is subject to a major challenge: the combination problem. This is roughly the question: how do the experiences of fundamental physical entities such as quarks and photons combine to yield the familiar sort of human conscious experience that we know and love.

The most influential formulation of the combination problem was given by William James in *The Principles of Psychology* (1895). In criticizing “mind-dust theory”, on which mental states are held to be compounds of elemental mental states, James made the following observations:

Where the elemental units are supposed to be feelings, the case is in no wise altered. Take a hundred of them, shuffle them and pack them as close together as you can (whatever that may mean); still each remains the same feeling it always was, shut in its own skin, windowless, ignorant of what the other feelings are and mean. There would be a hundred-and-first feeling there, if, when a group or series of such feelings

---

<sup>0</sup>Forthcoming in (Ludwig Jaskolla and Godehard Bruntrup, eds) *Panpsychism*, Oxford University Press. This article is based on my opening presentation to “Panpsychism on the Reef,” a workshop on the combination problem held on Lady Elliot Island in July 2012. Thanks to the audience there for useful discussion. The material on the structural mismatch problem benefited from discussion at the Oslo conference on panpsychism in August 2013. I am grateful to many philosophers for their responses to the first draft of this paper, both in conversation and in print. In some cases I have added some new discussion in response, mainly in footnotes. For written comments I am grateful to John Gregg and Tom McClelland.

were set up, a consciousness belonging to the group as such should emerge. And this 101st feeling would be a totally new fact; the 100 original feelings might, by a curious physical law, be a signal for its creation, when they came together; but they would have no substantial identity with it, nor it with them, and one could never deduce the one from the others, or (in any intelligible sense) say that they evolved it.

Take a sentence of a dozen words, and take twelve men and tell to each one word. Then stand the men in a row or jam them in a bunch, and let each think of his word as intently as he will; nowhere will there be a consciousness of the whole sentence. We talk of the ‘spirit of the age,’ and the ‘sentiment of the people,’ and in various ways we hypostatize ‘public opinion.’ But we know this to be symbolic speech, and never dream that the spirit, opinion, sentiment, etc., constitute a consciousness other than, and additional to, that of the several individuals whom the words ‘age,’ ‘people,’ or ‘public’ denote. The private minds do not agglomerate into a higher compound mind.

James is here arguing that experiences (feelings) do not aggregate into further experiences, and that minds do not aggregate into further minds. If this is right, any version of panpsychism that holds that microexperiences (experiences of microphysical entities) combine to yield macroexperiences (experiences of macroscopic entities such as humans) is in trouble.

In recent years, there has been a small groundswell of activity on panpsychism, and in particular there has been a small groundswell of activity on the combination problem. The problem was given its name by William Seager (1995) and was given an especially sharp formulation by Philip Goff (2009). Proposals for addressing it have been presented by Sam Coleman (2012, 2013, this volume), Goff (2009b, 2011, this volume), Gregg Rosenberg (2004, 2014), Seager (2010, this volume), and others.<sup>1</sup> It is fair to say that no proposed solution has yet gained much support, however.

This article is an attempt at a systematic treatment of the combination problem. I distinguish a number of aspects or versions of the problem. I discuss various ways in which the combination problem can be turned into an argument against panpsychism. I then try to systematically lay out the options for dealing with the combination problem, examining their advantages and disadvantages.

---

<sup>1</sup>Other recent work discussing the combination problem includes Basile 2012, Blamauer 2011, Dainton 2011, Gabora 2002, Goff 2006, Hunt 2011, Montero (this volume), Mørch 2014, Roelofs 2014, Shani 2010, Skrbina 2011, and Strawson 2006b.

A reasonable goal here is to either solve the combination problem or prove that it cannot be solved. I cannot say that I have achieved either of these objectives in this article as it stands, but I hope to at least clarify the issues enough to help others to make progress.

## 2 Terminology

First, some terminology. Most of this terminology is drawn from “Panpsychism and Panprotopsychism”, so the presentation here is much briefer than the presentation there. Interested readers may well find it useful to read the other article first, though the current article is self-contained in principle.

Microphysical properties and entities are the fundamental physical properties and entities characterized by a completed physics. Phenomenal properties are properties characterizing what it is like to be a conscious subject. Microphenomenal properties are the phenomenal properties of microphysical entities. Macrophenomenal properties are the phenomenal properties of other entities, such as humans. Microphenomenal and macrophenomenal truths are truths about the instantiation of these properties.

Constitutive panpsychism is the thesis that macrophenomenal truths are (wholly or partially) grounded in microphenomenal truths. Nonconstitutive panpsychism is the thesis that macrophenomenal truths are not grounded in microphenomenal truths. The most important form of nonconstitutive panpsychism is emergent panpsychism, on which macrophenomenal properties are strongly emergent from microphenomenal or microphysical properties, perhaps in virtue of fundamental laws connecting microphenomenal to macrophenomenal.

Russellian panpsychism is the thesis that microphenomenal properties are quiddities: the categorical bases of fundamental microphysical dispositions, or the properties that play fundamental microphysical roles. For example, the quiddity associated with mass is the property that plays the mass role (resisting acceleration, attracting other masses, and so on). Numerous philosophers have argued that the nature of quiddities is hidden from us. The Russellian panpsychist holds that quiddities are themselves phenomenal.

Perhaps the most important form of panpsychism is constitutive Russellian panpsychism, on which microphenomenal properties serve as quiddities and also serve to constitute macrophenomenal properties. I argue in “Panpsychism and Panprotopsychism” that this view is better-suited than any other form of panpsychism to deal with the problem of mental causation. On this view, microphenomenal properties are causally efficacious in virtue of their playing funda-

mental microphysical roles, and macrophenomenal properties are causally efficacious in virtue of being grounded in microphenomenal properties. By contrast, nonconstitutive and nonRussellian panpsychism have many of the same problems with mental causation as dualism.

Panprotopsychism is the thesis that fundamental physical entities have protophenomenal properties. Protophenomenal properties are special properties that are not themselves phenomenal (there is nothing it is like to have them) but that can collectively constitute phenomenal properties. To rule out standard forms of materialism from counting as panprotopsychism, these special properties must be (i) distinct from the structural/dispositional properties of microphysics and (ii) their constitutive relation to phenomenal properties must reflect an a priori entailment from protophenomenal to phenomenal truths.

Constitutive panprotopsychism is the thesis that macrophenomenal truths are grounded in truths about the protophenomenal properties of microphysical entities. Russellian panprotopsychism is the thesis that protophenomenal properties serve as quiddities. Constitutive Russellian panprotopsychism is perhaps the most important form of panprotopsychism, for the same reasons as in the case of constitutive Russellian panpsychism.

### **3 The many combination problems**

The combination problem for panpsychism is: how can microphenomenal properties combine to yield macrophenomenal properties? The combination problem for panprotopsychism is: how can protophenomenal properties combine to yield macrophenomenal properties? I will concentrate especially on the problem for panpsychism, but I will address both.

The combination problem can be broken down into at least three subproblems, reflecting three different aspects of phenomenal states: their subjective character (they are always had by a subject), their qualitative character (they involve distinctive qualities), and their structural character (they have a certain complex structure). These three aspects yield what we might call *the subject combination problem*, the *quality combination problem*, and the *structure combination problem*.

The subject combination problem is roughly: how do microsubjects combine to yield macrosubjects? Here microsubjects are microphysical subjects of experience, and macrosubjects are macroscopic subjects of experience such as ourselves.

An especially pressing aspect of the subject combination problem is the *subject-summing* problem. One can pose this problem by an extension of James' reasoning in the passage quoted earlier. Given 101 subjects, it seems that the existence of the first 100 does not necessitate the existence of

the 101st. More generally, given any group of subjects and any further subject, it seems possible in principle for the first group of subjects to exist without the further subject. If so, then no group of microsubjects necessitates the existence of a macrosubject.

The quality combination problem is roughly: how do microqualities combine to yield macroqualities? Here macroqualities are specific phenomenal qualities such as phenomenal redness (what it is like to see red), phenomenal greenness, and so on. It is natural to suppose that microexperience involves microqualities, which might be primitive analogs of macroqualities. How do these combine?

An especially pressing aspect of the quality combination problem is what we might call the *palette problem*.<sup>2</sup> There is a vast array of macroqualities, including many different phenomenal colors, shapes, sounds, smells, and tastes. There is presumably only a limited palette of microqualities. Especially if Russellian panpsychism is true, we can expect only a handful of microqualities, corresponding to the handful of fundamental microphysical properties. How can this limited palette of microqualities combine to yield the vast array of macroqualities?

The structure combination problem is roughly: how does microexperiential structure (and microphysical structure) combine to yield macroexperiential structure? Our macroexperience has a rich structure, involving the complex spatial structure of visual and auditory fields, a division into many different modalities, and so on. How can the structure in microexperience and microstructure yield this rich structure?

An especially pressing aspect of the structure combination problem is the *structural mismatch* problem. Macrophysical structure (in the brain, say) seems entirely different from the macrophenomenal structure we experience. Microexperiences presumably have structure closely corresponding to microphysical structure (this is especially clear on a Russellian view), and we might expect a combination of them to yield something akin to macrophysical structure. How do these combine to yield macrophenomenal structure instead?

There are a few other aspects of the combination problem, corresponding to different aspects of macroexperience that need explaining. There is the *unity problem*: how do microexperiences come together to yield a unified consciousness? There is the *boundary problem* (Rosenberg 1998): how do microexperiences come together to yield a bounded consciousness? There is the *awareness problem*: how do microexperiences come together to yield awareness of qualities? And there is the *grain problem* (Maxwell 1979; Lockwood 1993): how do microexperiences come together

---

<sup>2</sup>Dainton (2011) calls this problem the “derivation problem”.

to yield homogeneous macroexperiences, such as a homogeneous experience of red, instead of an enormous jagged array of distinct qualities? Some of these problems might be assimilated to earlier problems (the first three plausibly involve aspects of subjective character, the last involves an aspect of qualitative character, and all involve aspects of structure), but it is useful to have them on the table explicitly.

It is common for a proposed solution to the combination problem to address only one of these problems: most often the subject combination problem and occasionally the quality combination problem. It should be stressed that a satisfactory solution to the combination problem must address all of these problems. This raises the bar for a solution, as it is far from clear that any single proposal can solve all the problems at once. One might appeal to separate proposals for solving the problems one at a time, but then it is far from clear that these proposals will be compatible with each other. At the very least, any proposed solution to the combination problem should indicate which problems it is addressing, and which problems it is not.

The formulation of the problems above is misleading in one respect. I have typically said “how do microexperiences come together to yield X”, or perhaps “how do microsubjects” or “how do microqualities”. However, constitutive panpsychism is not committed to the claim that macroexperience is *wholly* grounded in microexperience. It could be partly grounded in causal or structural relations among the microexperiences, or in other microphysical properties, or even in other quiddities if there are non-phenomenal quiddities as well. We can put all this by saying that constitutive panpsychism requires macroexperiences to be wholly grounded in microexperiences and microphysics, where microphysics is understood broadly to include all of the above. The formulations of the relevant problems can then all take the form “How do microexperiences and microphysics come together to yield X?”. With the problems understood this way, the panpsychist has more resources to play with, but the problems still seem very difficult to solve.

There are analogous versions of all of the problems for panprotopsychism. We need only replace the appeal to microexperience with an appeal to protoexperience (the instantiation of proto-phenomenal properties), yielding questions of the form: “How do protoexperiences come together to yield X?”, or “How do protoexperiences and microphysics come together to yield X?”.

The structure and quality combination problems seem just as hard in this guise. The subject combination problem will take a different form, one that perhaps makes it slightly easier. Although microexperiences presumably have subjects, protoexperiences need not. Panprotopsychism therefore need not appeal to microsubjects, and need not require subjects to combine into other subjects. Still, there remains a substantial challenge in explaining how non-subjects of experience can com-

bine to yield subjects of experience.

More generally, panprotopsychism faces a version of the combination problem that does not arise for panpsychism: how can nonexperiences constitute experiences? Sometimes it is flatly asserted that this is impossible, or it is suggested that it is a general gap between the nonexperiential and the experiential that underlies and explains the gap between the physical and the experiential. I do not think that this is obviously correct: I think one can point to special features of the physical that explain the latter gap (the structural nature of physical truths, for example), and I have not seen any argument for a general nonexperiential-experiential gap that is as powerful as the arguments for an physical-experiential gap. Still, there is at least a significant challenge for the panprotopsychist here.

I conclude that both panpsychism and panprotopsychism suffer from serious combination problems.

## **4 Turning the combination problem into an argument**

Can the combination problem be proved unsolvable? That requires, in effect, turning the challenges posed by the combination problem into a conclusive argument against pan(proto)psychism, or at least against constitutive Russellian pan(proto)psychism. Of course conclusive arguments are hard to come by in philosophy, but we can at least examine the arguments that are available.

### **4.1 The anti-aggregation argument**

One of James' central arguments against panpsychism in *The Principles of Psychology* appeals to the general thesis that aggregates do not really exist: a view sometimes called nihilism about composition. More precisely, James holds that aggregates do not have objective existence, but exist only for observers who perceive them as such. He writes:

In other words, no possible number of entities (call them as you like, whether forces, material particles, or mental elements) can sum themselves together. Each remains, in the sum, what it always was; and the sum itself exists only for a bystander who happens to overlook the units and to [p.159] apprehend the sum as such; or else it exists in the shape of some other effect on an entity external to the sum itself. Let it not be objected that H<sub>2</sub> and O combine of themselves into 'water,' and thenceforward exhibit new properties. They do not. The 'water' is just the old atoms in the new

position, H-O-H; the ‘new properties’ are just their combined effects, when in this position, upon external media, such as our sense-organs and the various reagents on which water may exert its properties and be known.

We might try turning this into an argument as follows:

- (1) If constitutive panpsychism is true, human consciousness is an aggregate.
  - (2) Aggregates do not objectively exist.
  - (3) Human consciousness objectively exists.
- 
- (4) Constitutive panpsychism is false.

The key premise is premise (2). Its support is the claim that aggregates exist only for observers, or only in virtue of their effects. James does not give much support for these claims, however, and they are easy to reject. A more orthodox view holds that aggregate entities such as molecules exist independently of observers and independently of their effects. Of course James’ nihilism about composite objects is not indefensible. Still, nihilist theses of this sort are so widely rejected that they do not have much dialectical force in an argument against panpsychism.<sup>3</sup>

There is perhaps some intuitive force to the idea that consciousness has a higher and purer degree of existence than tables and molecules. A related argument (consistent with the framework of my “Ontological Anti-Realism”) proceeds from the claim holds that conscious subjects exist *determinately* whereas aggregates do not. This argument does not require nihilism and arguably applies more plausibly to conscious subjects than to rocks. Still, the premise that aggregates do not determinately exist is highly controversial, so the dialectical force of the argument remains limited.

## 4.2 The subject-summing argument

The subject-summing argument is suggested by James’ argument, quoted at the start of the paper, against combination of feelings and minds. We can formalize an argument roughly as follows. As with all the arguments I present in this section and the next, this formalization largely follows the way that closely related arguments are presented in Goff (2009).

---

<sup>3</sup>Shani (2010) gives a thorough discussion of James’ argument here.



(1) If constitutive panpsychism is true, the existence of a number of microsubjects with certain experiences necessitates the existence of a distinct macrosubject.

(2) It is never the case that the existence of a number of subjects with certain experiences necessitates the existence of a distinct subject.

---

(3) Constitutive panpsychism is false.

Strictly speaking, premises (1) and (2) should allow arbitrary microphysical truths to be conjoined with the truths about the subjects, but the simple version conveys the main point. Premise (2) is the key premise. An intuitive case for either version of it can be made along the lines of the quote from James at the start of the paper. One can also support it using a further conceivability argument.

(1) For any group of subjects (with certain experiences), it is conceivable that those subjects exist (with their experiences) and no other subjects exist.

(2) For any group of subjects, if it is conceivable that that those subjects exist (with their experiences) and no other subjects exist, then this is possible.

---

(3) For any group of subjects (with certain experiences), it is possible that the subjects in  $S$  exist (with their experiences) and no other subjects exist

Premise (1) has a reasonable degree of intuitive support. Even when adjusted to allow arbitrary microphysical truths to be conjoined with the existence of the subjects in  $S$ , it retains considerable support. Premise (2) is an instance of a general conceivability/possibility claim. Of course conceivability/possibility claims can be rejected, but not without incurring substantial costs, and panpsychists who have rejected physicalism in part on the basis of conceivability arguments are not in a good position to do so. So this argument poses a significant challenge to the constitutive panpsychist. I will examine options for answering it later.

### 4.3 The conceivability argument

The preceding considerations suggest a more general conceivability argument against constitutive panpsychism, inspired by the conceivability argument against physicalism. Here  $PP$  is a

conjunction of all microphysical and microphenomenal truths about the universe, while  $Q$  is a macrophenomenal truth, such as ‘Some macroscopic entity is conscious’.

- (1)  $PP \& \neg Q$  is conceivable.
- (2) If  $PP \& \neg Q$  is conceivable, it is metaphysically possible.
- (3) If  $PP \& \neg Q$  is metaphysically possible, constitutive panpsychism is false.

---

(4) Constitutive panpsychism is false.

Here premises (2) and (3) parallel corresponding premises in the familiar conceivability argument against physicalism (e.g. Chalmers 2009b). The distinct premise is (1). This premise in effect asserts the conceivability of a *panpsychist zombie world*: a world in which microphysics and microexperience is just as it is in our world, but in which no macroscopic entity is conscious. Such a world is populated by panpsychist zombies, which are microphysical and microphenomenal duplicates of us without consciousness.

Why believe premise (1)? One might think it has a certain intuitive force, just as does the corresponding premise about the conceivability of microphysical-duplicate zombies. However, one can also support it by appealing to the first premise of the conceivability argument in the last section. If we appeal to the modified version of that premise, saying that for any group of conscious subjects and any microphysical truths, it is conceivable that the microphysical truths obtain and the subjects in that group exist without any other subjects, then premise (1) follows.

One might also support premise (1) in other ways. One could use considerations about the quality combination problem to support it, for example arguing that one can conceive of arbitrary microqualities without distinct macroqualities. One could also use considerations about the structure combination problem to support it, arguing that one can conceive of microphenomenal and microphysical structure without distinct macrophenomenal structure. Many of these principles will also generate direct arguments against panpsychism in their own right, but it is useful to have the argument above in the arsenal.

There is also a conceivability argument against panprotopsychism, which replaces  $PP$  in the argument above by  $PPP$ , the conjunction of protophenomenal and microphysical truths. The key premise (1) will now say that  $PPP \& \neg Q$  is conceivable. Why believe this? One might again think it has intuitive support, though this is far from clear given that we have so weak a conception of what protophenomenal properties are like. Alternatively, it might gain support from a thesis

holding that for any nonexperiential truth  $N$  and any experiential truth  $E$ ,  $N \& \neg E$  is conceivable. Once again, such a thesis might generate a direct argument against panprotopsyichism in its own right, but the argument form above helps clarify the territory.

#### 4.4 The knowledge argument

Having considered the conceivability argument, it is natural to consider a knowledge argument. We can suppose that inside her black-and-white room, Mary is told all the microphysical facts, and also learns all the microphenomenal facts: she learns what it is like to be a quark, a photon, and so on. Perhaps this is accomplished by giving her versions of those experiences, or by somehow enabling her to imagine them. One might think that in this situation, Mary would still be unable to know what it is like to see red, even given arbitrary a priori reasoning. If so, one could mount an argument as follows. Here  $PP$  is as before and  $Q$  is a truth about what it is like to see red, and “deducible” means “inferrable by a priori reasoning alone”.

- (1)  $Q$  is not deducible from  $PP$ .
- (2) If  $Q$  is not deducible from  $PP$ ,  $Q$  is not necessitated by  $PP$ .
- (3) If  $Q$  is not necessitated by  $PP$ , constitutive panpsychism is false.

---

- (4) Constitutive panpsychism is false.

There is also a corresponding argument against panprotopsyichism that replaces microphenomenal facts by protophenomenal facts, and replaces  $PP$  by  $PPP$  above. One might likewise think it intuitive that knowledge of all the protophenomenal facts would not help Mary to know what it is like to see red.

I think that these arguments are highly inconclusive, largely because we know so little about what microphenomenal or protophenomenal properties are like. Perhaps once we grasped them, we would understand their connection to experiences of red and to other experiences. Certainly there does not seem to be a general case for premise (1) here that is nearly as strong as the case for the premise involving microphysical truths alone. Still, perhaps such a case might be mounted.

#### 4.5 The palette argument

I turn next to an argument associated with the quality combination problem, inspired by the palette problem discussed earlier.

(1) If constitutive panpsychism is correct, macrophenomenal qualities are constituted by microphenomenal qualities.

(2) If Russellian panpsychism is correct, there are only a few microphenomenal qualities.

(3) Macrophenomenal qualities are too diverse to be constituted by a few microphenomenal qualities.

---

(4) Constitutive Russellian panpsychism is incorrect.

Where the previous arguments were arguments against constitutive panpsychism in both Russellian and non-Russellian varieties, this one is an argument against only the former. Russellian panpsychism requires that microphenomenal properties are all directly associated with a fundamental physical property, and there appear to be only a few of these. Non-Russellian panpsychism, by contrast, can escape the argument by allowing that there is a diverse array of microphenomenal qualities.

The case for the key premise (3) is intuitive and inconclusive as it stands: perhaps we might find a small set of deep underlying qualities with sufficient generality to generate all phenomenal qualities, just as we have done for physical qualities. But this is at least an argument that panpsychists need to address.

#### **4.6 The revelation argument**

The next argument is also loosely associated with the quality combination problem, and is especially closely associated with the grain problem discussed earlier. Versions of this argument are discussed by Lockwood (1993) and Goff (2006).

(1) The nature of consciousness is revealed to us in introspection.

(2) If constitutive panpsychism is correct, consciousness is constituted by a vast array of microexperiences.

(3) Whatever constitutes consciousness is part of its nature.

(4) A vast array of microexperiences is not revealed to us in introspection.

---

(5) Constitutive panpsychism is incorrect.

Premise (1) is not compulsory, and most materialists will deny it. But the premise nevertheless has a certain intuitive plausibility, and some theorists invoke something like it to argue against materialists. For panpsychists who argue in this way, it is an uncomfortable premise to deny. Premises (2) and (4) are also hard to deny.

Perhaps the best way to respond to this argument is to deny premise (3). One can distinguish the nature of a phenomenal property from the grounds (or realizers or constituters) of an instance of that property. It is a familiar point that a single property can be multiply realized by different grounds in different instances, and it is not clear why the same should not also apply to phenomenal properties. It is then coherent to hold that the nature of a phenomenal property is revealed by introspection although the grounds of a specific instance are not.

#### **4.7 The structural mismatch argument**

This argument is inspired by the structural mismatch problem discussed earlier: macrophenomenal structure (of consciousness) seems quite different from macrophysical structure (of the brain, say) where constitutive Russellian panpsychism would seem to require that the structures be the same. It is also closely related to the grain problem, which is used (for example by Maxwell and Stoljar) to raise a version of the structural mismatch problem

We can understand microphysical structure and macrophysical structure as the quasi-mathematical structure of microphysical and macrophysical entities as characterized by physics. Macrophenomenal structure is the structure we find within our phenomenology. In both cases, structure includes both internal structure (the internal geometrical structure of a complex physical entity, the internal structure of a visual field) as well as what we might external structure: the structure of spaces within which properties are embedded (the scalar structure of mass, the three-dimensional structure of color space).

The structural mismatch argument can be put in the form of an apparently inconsistent tetrad:

- (1) Microphenomenal structure is isomorphic to microphysical structure
- (2) Microphenomenal structure constitutes macrophenomenal structure.
- (3) Microphysical structure constitutes macrophysical structure.
- (4) Macrophenomenal structure is not isomorphic to macrophysical structure.

Here (1) is an apparent commitment of Russellian panpsychism, (2) is an apparent commitment of constitutive panpsychism, and (3) is a widely accepted view of the physical. (4) reflects

the plausible datum of mismatch between the structure of consciousness and the structure of the brain. When combined with the additional premise saying that (1)-(4) are inconsistent, it follows that constitutive Russellian panpsychism is false.

A corresponding argument against panprotopsyichism replaces “microphenomenal” by “protophenomenal” in premises (1) and (2). These premises are then apparent commitments of Russellian and constitutive panprotopsyichism respectively, so that the inconsistency of the premises yields an argument against constitutive Russellian panprotopsyichism.

Although the structural mismatch argument has received relatively little attention to date, I think it is one of the more powerful arguments against constitutive Russellian versions of panpsychism and panprotopsyichism. There are various ways to respond to the argument, but doing so is not at all trivial. I consider the argument at some length later in this article.

## 5 Noncombinatorial responses

The most obvious sort of panpsychist response to the combination problem is a *combinatorial* response: show how microexperiences can constitutively combine to yield macroexperiences. But there are also *noncombinatorial* responses, which deny that microexperiences constitutively combine to yield macroexperiences.

The most obvious sort of noncombinatorial response is emergent panpsychism, which holds that macroexperiences are strongly emergent from microexperiences and are not constituted by them. This view rejects constitutive panpsychism, so it does not need to give an account of mental combination.

Another noncombinatorial response is *identity panpsychism*, on which macroexperiences are identical to microexperiences. On this view, macroexperiences are already present at the fundamental level and no combination is required. Given that microexperiences constitute themselves, this view is nevertheless a form of constitutive panpsychism.

A third noncombinatorial response is *autonomous panpsychism*, which holds that macroexperiences are autonomous from microexperiences, in that they are neither constituted by, emergent from, nor identical to microexperiences. On one version of this view, microexperiences are emergent from or constituted by macroexperiences. On another version, microexperiences and macroexperiences are both autonomous, with neither depending on the other.

These three noncombinatorial responses contrast with the more familiar *combinatorial panpsychism*, on which microexperiences collectively constitute macroexperiences. It is worth noting that

all forms of panprotopsychism are combinatorial: by definition, protophenomenal properties are distinct from but can collectively constitute phenomenal properties.

This taxonomy divides panpsychist responses to the combination problem into four classes: emergent panpsychism, autonomous panpsychism, identity panpsychism, and combinatorial panpsychism. The first three are noncombinatorial responses, while the third is a combinatorial response. The last two are forms of constitutive panpsychism, while the first two are forms of nonconstitutive panpsychism.

Each of these four broad classes subsumes various specific sorts of response in turn. In this section I discuss the noncombinatorial responses: emergent panpsychism, autonomous panpsychism and identity panpsychism.

### **5.1 Emergent panpsychism**

Emergent panpsychism holds that macroexperiences are not grounded in microexperiences, but instead are strongly emergent from microexperiences, from microphysics, or from both. Strong emergence involves the emergence of ontologically novel entities that are not grounded in the base entities. On a common conception of strong emergence, the base entities do not metaphysically necessitate the emergent entities, but instead they are connected by contingent laws of nature. On this conception of emergent panpsychism, there will be contingent laws of nature connecting microexperience (or microphysics) to macroexperience.

Emergent panpsychism has the great advantage of avoiding the combination problem. Strongly emergent entities and properties are best construed as fundamental entities and properties, not grounded in the base entities or in other entities. As such, no combination is required (except, perhaps, insofar as we construe the laws connecting microexperience with macroexperience as laws of combination). On this view, macrosubjects are fundamental entities, just as they are according to substance dualism. This allows emergent panpsychism to avoid the combination problem just as substance dualism does.

At the same time, emergent panpsychism shares many of the disadvantages of substance dualism. It suffers from problems of economy, postulating many more fundamental entities in the world. And perhaps more important, it suffers from the problems of mental causation. Because macroexperience is not grounded in microphysics or microexperience, it cannot inherit the causal relevance of either. Given that microphysics is causally closed, it is hard to see how macroexperience can have any causal effects on it. Like substance dualism, emergent panpsychism seems to

face an unattractive choice between epiphenomenalism, interactionism, and underdetermination.

Of course this does not mean that emergent panpsychism is not true. It may be that it has other advantages over substance dualism, for example with respect to continuity and elegance. It certainly has the advantage of avoiding the combination problem! But for those (like me) who are interested in panpsychism in large part because it promises to avoid the problems of mental causation, emergent panpsychism seems to sacrifice this motivation.

Many solutions to the combination problem that have been put forward turn out on close examination to be forms of emergent panpsychism. For example, Gregg Rosenberg (2004) invokes ontologically primitive “high-level individuals” that emerge from lower-level individuals. Liane Gabora (2002) invokes fundamental principles for “amplifying phenomenal information”, in virtue of which macroexperience strongly emerges from microexperience. Giulio Tononi’s integrated information theory (2008), which puts forward a principle connecting degrees of integrated information with states of consciousness, can also be construed as a form of emergent panpsychism. If we see Tononi’s principle as a fundamental law of nature, then it appears that macroexperiences are strongly emergent from certain physical configurations.

Many of these theorists do not deal directly with the problem of mental causation. Rosenberg is an exception: he deals with mental causation by allowing high-level individuals to exert a small amount of downward causation through interaction with the underlying entities. I think that once it becomes clear that these solutions are subject to the same worries about mental causation as substance dualism, they lose some of their initial attractions. Again, this is not to say that these theories are false. But it does give us motivation to look elsewhere.<sup>4</sup>

---

<sup>4</sup>Hedda Hassel Mørch (2014) defends emergent panpsychism by holding that (i) emergent causal relations can be intelligible rather than brute, (ii) macroexperiences have metaphysical priority over the microexperiential parts from which they emerge, and (iii) macroexperiences are the intrinsic natures of certain macroscopic physical systems that have metaphysical priority over their microphysical parts. Mental causation is handled by the observation that macroexperiences are more fundamentally efficacious than their microphysical parts. Challenges for this view include understanding how macroscopic entities and properties can be metaphysically prior to the microscopic entities and properties that cause them, and understanding how and whether intrinsically identical microscopic entities and properties will be causally efficacious or not depending on their macroscopic surrounds. One could handle these challenges by taking the view to be a version of the “combinatorial infusion” view (discussed later) with fundamental fused entities in the physics, but the various challenges for that view must then be met.



## 5.2 Autonomous panpsychism

Like emergent panpsychism, autonomous panpsychism denies that macroexperiences are grounded in microexperiences. Unlike emergent panpsychism, it denies that macroexperiences are even strongly emergent from microexperiences (or from microexperiences and microphysics). In effect, emergent panpsychism retains a sort of dependence of macroexperience on microexperience, if a dependence weaker than grounding or constitution (an asymmetrical nomological dependence, perhaps). Autonomous panpsychism denies even this weak sort of dependence.<sup>5</sup> As a result, it is not easy to square the view with a contemporary worldview on which everything depends at least weakly on what is going on in physics, but the view is worthy of some attention.

One version of autonomous panpsychism says that microexperiences are grounded in (or constituted by) macroexperiences, so that macroexperiences are fundamental and microexperiences are derivative. On one version of this view, human-level experiences are fundamental, as on certain forms of idealism. On another version, universe-level experiences (experience of the whole universe as a subject) are fundamental: a sort of cosmopsychism. These views have to deal with a reverse version of the combination problem, which we might call the decomposition problem. How does macroexperience give rise to microexperience? For example, how does a single subject give rise to multiple dependent subjects? How do macroqualities yield microqualities, and how does macroexperiential structure yield microexperiential structure? These problems seem just as hard as the original combination problem.

Another version of autonomous panpsychism says that microexperiences are strongly emergent from macroexperiences, while a third version says that neither microexperiences nor macroexperiences depend on the other. On these views (as on emergent panpsychism), both microexperiences and macroexperiences are metaphysically fundamental. As with emergent panpsychism and substance dualism, these views avoid the combination and decomposition problems (at least in their hardest forms), but they face the problem of mental causation. On these views microexperiences and macroexperiences are both fundamental, so it appears that they will compete for causal relevance. Given the causal closure of the microphysical, it appears that we have a familiar choice between epiphenomenalism and overdetermination at the macroexperiential level. Again, this does not show that autonomous panpsychism is false, but it gives some motivation for looking

---

<sup>5</sup>In forthcoming work, Jennifer McWeeny argues that the seventeenth-century philosopher Margaret Cavendish was a sort of autonomous panpsychist, holding that everything in the universe is conscious and that consciousness at one level does not depend on consciousness at other levels.

at alternative solutions.

### 5.3 Identity panpsychism

Identity panpsychism holds that macroexperiences that are identical to microexperiences: experiences had by fundamental physical entities. This view requires that macrosubjects are themselves microsubjects, or fundamental physical entities. The view may sound unpromising at first, but versions of it are worth exploring.

The version of identity panpsychism that first comes to mind is what we might call the *dominant monad* view, by analogy to Leibniz's view on which are identical to a single localized monad. On this view, the subject of our experiences is a single localized fundamental entity: perhaps a single quark somewhere in our brain. The microexperiences of this quark are precisely our macroexperiences. There are obvious worries here about this quark's stability (what happens when it disappears?) and about its causal role (how could its properties play the rich causal role that macroexperiences seem to play?).

Even harder problems arise when the view is combined with Russellian panpsychism, on which microphenomenal properties correspond directly to microphysical properties. For a start, the quark is presumably microphysically like other quarks, so it will also be microphenomenally like those quarks, yielding a vast manifold of subjects of experience just like me throughout the brain and throughout the universe. And given the simplicity of the microphysical structure of a quark, it is hard to see how the corresponding microphenomenology could have anything like the complexity of our macroexperience. So unless this view is combined with serious revisions to physics, it is probably best put aside.

Other versions of identity panpsychism are *holistic* in that they invoke fundamental physical entities that are not atomic or localized. One such view combines identity panpsychism with the monistic view that the universe itself is the most fundamental physical entity. The result is *identity cosmopsychism*, on which the whole universe is conscious and on which we are identical to it. (Some idealist views in both Eastern and Western traditions appear to say something like this.) Obvious worries for this view are that it seems to entail that there is only one conscious subject, and that each of us is identical to each other and has the same experiences. There is also a structural mismatch worry: it is hard to see how the universe's experiences (especially given a Russellian views on which these correspond to the universe's physical properties) should have anything like the localized idiosyncratic structure of my experiences. Perhaps there are sophisticated versions of

this view on which a single universal consciousness is differentiated into multiple strands of mid-level macroconsciousness, where much of the universal consciousness is somehow hidden from each of us. Still, this seems to move us away from identity cosmopsychism toward an autonomous cosmopsychist view in which each of us is a distinct constituent of a universal consciousness. As before, the resulting decomposition problem seems just as hard as the combination problem.

Perhaps the most important version of identity panpsychism is *quantum holism*. This view starts from the insight that on the most common understandings of quantum mechanics, the fundamental entities need not be localized entities such as particles. Multiple particles can get entangled with each other, and when this happens it is the whole entangled system that is treated as fundamental and that has fundamental quantum-mechanical properties (such as wave functions) ascribed to it. A panpsychist might speculate that such an entangled system, perhaps at the level of the brain or one of its subsystems, has microphenomenal properties. On the quantum holism version of identity panpsychism, macrosubjects such as ourselves are identical to these fundamental holistic entities, and our macrophenomenal properties are identical to its microphenomenal properties.

This view has more attractions than the earlier views, but there are also worries. Some worries are empirical: it does not seem that there is the sort of stable brain-level entanglement that would be needed for this view to work. Some related worries are theoretical: On some interpretations of quantum mechanics the locus of entanglement is the whole universe (leading us back to cosmopsychism), on others there is no entanglement at all, and on still others there are regular collapses that tend to destroy this sort of entanglement. But perhaps the biggest worry is once again a structural mismatch worry. The structure of the quantum state of brain-level systems is quite different from the structure of our experience. Given a Russellian view on which microphenomenal properties correspond directly to the fundamental microphysical properties of these entangled systems, it is hard to see how they could have the familiar structure of our macroexperience.

The identity panpsychist (of all three sorts) might try to remove some of these worries by rejecting Russellian panpsychism, so that microphenomenal properties are less closely tied to microphysical structure. The cost of this move is that it becomes much less clear how these phenomenal properties can play a causal role. On the face of it they will be either epiphenomenal, or they will make a difference to physics. The latter view will in effect require a radically revised physics with something akin to our macrophenomenal structure present at the basic level. Then phenomenal properties will in effect be playing the role of quiddities within this revised physics, and the resulting view will be a sort of revisionary Russellian identity panpsychism.

The overall moral is that it is difficult for the identity panpsychist to avoid epiphenomenalism on one hand or radical revisions in physics on the other. Still, at least the quantum holist version of the view deserves close examination.

## 6 Combinatorial responses

The most important class of responses to the combination problems are combinatorial responses, on which microexperiences (or protoexperiences) collectively constitute macroexperiences. Here numerous strategies are available. I will start by considering strategies for dealing with the subject combination problem, and will then consider strategies for dealing with the other problems.

### 6.1 Deflate the subject

Any combinatorial version of panpsychism or panprotopsychism must be at least somewhat deflationary about subjects of experience. If subjects were metaphysically primitive entities, they could not be constituted by more basic entities, and combinatorial views would be ruled out. So these views must deny that subjects are metaphysically primitive entities. Indeed, proponents of these views might argue that the subject-summing argument is generated by a tacit background presupposition that subjects are metaphysically primitive entities. If this is right, then replacing this presupposition with a more adequate view of subjects might hold the key to solving the subject combination problem.

An extreme form of deflationism about subjects is *eliminativism*: the view that there are no subjects of experience. If this view is correct, then there are no macrosubjects and the subject combination problem does not need to be addressed. Many of the great neutral monists (themselves panprotopsychists), such as Mach, James, and Russell at least flirted with this sort of eliminativism. Sometimes this view came down to denying a metaphysically primitive subject (as when Mach rejects an “ego” with “real unity” and James rejects a “soul”), but sometimes the view seems to take the more radical form of rejecting subjects altogether, as Russell does in *The Analysis of Matter* and James does in his work on radical empiricism.

Wholesale eliminativism about subjects is not easy to stomach, especially for someone who is serious about phenomenal properties. These properties are defined as those characterizing what it is like to be a subject. And however they are defined, as properties they presumably need bearers, which might then be taken to be subjects. So wholesale eliminativism about subjects may seem to

require eliminativism about phenomenal properties, or at least a reconception of them as properties of quite different entities.

Furthermore, eliminativism does not really remove all aspects of the subject combination problem. Presumably even an eliminativist will still acknowledge that experiences come in bundles or streams of some sort, so that the experiences previously taken to be mine share a bundle or stream with experiences previously taken to be yours. But now the problem can be reconceived as the bundle combination problem, or the stream combination problem: how can a number of distinct streams add up to a new single stream? It is not obvious that this problem is much easier than the original problem. Perhaps the eliminativist can also deny or deflate the bundling relation, but now the view is taking on even more significant costs.

Less extreme views hold that there are subjects while denying that they are metaphysically primitive: perhaps they are composite entities, or they are derivative entities in some other sense. This view has a number of attractions, and can be independently motivated by puzzle cases involving personal identity over time. Still, this sort of deflationism does not make the subject combination problem go away. We still need an account of how a derivative subject of experience can be constituted by microsubjects, or by microphenomenal/protophenomenal properties along with microphysics. Such a positive account is not easy to find, but I will consider some options in what follows.

## **6.2 Combinatorial infusion**

An idea that is sometimes mooted is that low-level subjects “merge” or “blend” or “fuse” to yield higher-level subjects. After the merging, the low-level subjects no longer exist in their own right. Only the higher-level subject exists. Seager (2010) calls this “combinatorial infusion”, on which a combined mental state “supersedes” the original mental states.

Many questions could be raised about this view, but a basic question is the following: is the relation between the original subjects and the merged subject a synchronic or a diachronic relation? If it is a synchronic relation, then presumably the low-level and high-level subjects exist at the same time, and we have lost the distinctive aspect of this view whereby the high-level subject supersedes the low-level subject. This version of the view will be faced with the original worries about how a number of subjects could ever synchronically constitute another subject.

Presumably the merging relation is diachronic, then. If so, it is hard to see how it can be a constitutive relation. Diachronic relations are naturally understood to be contingent causal relations,

not constitutive relations. Perhaps two subjects at an earlier time can nomologically necessitate the existence of a subject at a later time, but it is hard to see how they can metaphysically necessitate or constitute it. But constitution is what we need for a combinatorial solution to the combination problem.

If we examine the synchronic and constitutive structure of this view, it appears to be a form of noncombinatorial panpsychism. At the later time, there is a macrosubject and macroexperiences that are not constituted by microsubjects and microexperiences that exist at that time. So it appears that this macrosubject is itself fundamental. Either we have a version of emergent panpsychism, perhaps with this subject depending nomologically on underlying physical states, or we have a form of identity panpsychism, where this subject corresponds to a fundamental physical state.

One can bring out the point by asking how the view works as a form of Russellian panpsychism. Here the microsubjects and the microphenomenal properties must correspond directly to fundamental microphysical entities and their microphysical properties. So when a number of microsubjects go out of existence and are replaced by a “merged” subject, a number of microphysical entities presumably also go out of existence, replaced by a “merged” entity. This does not happen in classical physics, but it can happen in quantum physics. As Seager notes, when two particles become entangled, there is a sense in which neither exists any longer as a fundamental entity: instead they have “merged” into a fundamental entangled entity, of which the original particles are at best aspects.

The Russellian panpsychist could exploit this quantum-mechanical merging for their purposes, but the resulting position is a familiar one. It is a version of the quantum holism discussed under identity panpsychism in the previous section. It has the advantages and disadvantages discussed there (notably the worries about stability of entanglement and about structural mismatch), but it is not really a distinct view. Where constitutive relations are concerned, it is a form of identity panpsychism rather than combinatorial panpsychism.<sup>6</sup>

---

<sup>6</sup>Seager (this volume) suggests that his combinatorial infusion view can avoid various versions of the combination problem by appealing to laws of combinatorial infusion, which are fundamental laws of nature akin to laws of physics. On a constitutive Russellian panpsychist position, it is natural to hold that the only fundamental mental laws will be mental “realizations” of the fundamental laws of physics. Where physical properties are realized by mental quiddities, then laws connecting those properties will be realized by isomorphic laws connecting the corresponding quiddities. On this picture, any laws of combinatorial mental infusion must be realizers of a corresponding law of infusion in the fundamental physics, and the infused mental entity will realize an infused entity in the fundamental physics (a holistic quantum system, perhaps). This clearly leads to a form of identity panpsychism along the lines in the text, rather than a form of combinatorial panpsychism.

The challenge for this view is making the case that physics really contains infusion laws that yield infusions of the requisite character and complexity. As before, I think that the quantum holism version of the theory is the version most worth taking seriously, but it is not clear that the problems for that view can be overcome. One might try to find another source of infusion in physics, but I suspect that the worries that apply to quantum holism will probably still apply here.

As before, the merging theorist might reject the constitutive Russellian constraints, so that mental merging need not correspond to physical merging, but only at cost of raising serious worries about mental causation. For example, one could also understand laws of combinatorial infusion as “bridging” laws governing how multiple microsubjects combine to yield macrosubjects, but then the resulting picture appears to be a form of emergent panpsychism.

The overall upshot is that combinatorial infusion is best understood as a version of identity panpsychism or emergent panpsychism (with the associated problems), and not as a version of combinatorial panpsychism.<sup>7</sup>

### 6.3 Phenomenal bonding

Another suggestion (Goff 2009, this volume) is that microsubjects constitute macrosubjects in virtue of certain phenomenal *relations* between the microsubjects: phenomenal bonding relations. On this view, the subject-summing argument is generated in part by thinking of microsubjects as being merely related spatiotemporally or causally. Once we acknowledge distinctively phenomenal relations between microsubjects and their phenomenal states, we can see how all this might constitute a macrosubject and macrophenomenal states.

An immediate worry question is how there can be room for a phenomenal bonding relation, at least given a Russellian version of panpsychism. But there is an immediate answer. Microphysics postulates fundamental monadic properties such as mass and charge, but it also postulates fundamental relations such as spatiotemporal relations. Just as with mass and charge, physics seems

---

<sup>7</sup>In conversation, Tom Nagel has suggested a panprotopsychoist version of the infusion view, on which protophenomenal properties yield experiencing subjects as follows. If a fundamental physical entity is sufficiently isolated, its protophenomenal character determines an individual subject. If it is in the right kind of complex system, it instead contributes to determining a more complex (merged or infused) subject necessitated by the system as a whole. This view is in some ways reminiscent of Tononi’s integrated information theory, whose exclusion postulate says roughly that a system is conscious iff it is not part of (and does not contain) a system with a higher degree of integrated information. Both views seem to have the counterintuitive consequence that consciousness is extrinsic: intrinsically identical physical systems (with the same fundamental physical and protophenomenal properties) might be conscious or non-conscious depending on the surrounding context.

to characterize the mathematical structure of these relations but not their categorical nature. So just as monadic properties can have monadic quiddities underlying them as their categorical bases, relational properties might have relational quiddities underlying them as their categorical bases. It is then not out of the question that a certain phenomenal relation could serve as the quiddity underlying spatiotemporal relations.

A related idea (along the lines of Gregg Rosenberg's "carrier hypothesis" about causation in *A Place for Consciousness*) is that causation (or perhaps nomic necessitation) is a fundamental relation that has a phenomenal relation as an underlying quiddity. This version of the view also fits well with a Russellian phenomenal bonding theory. One could also invoke non-Russellian phenomenal bonding theories, but as always these will have trouble accommodating the causal relevance of phenomenal bonding and therefore of macrosubjects and macroexperience.

The biggest question for any phenomenal bonding view is: what is the phenomenal bonding relation? And how could any phenomenal relation holding between distinct subjects (or between phenomenal states of distinct subjects) suffice for the constitution of a wholly new subject?

A natural candidate here is the co-consciousness relation: a relation such that whenever it relates two phenomenal states, they are experienced jointly. When this relation holds among the states of distinct microsubjects, those states will be experienced jointly by a new subject.

One question for this view and for other phenomenal bonding views is whether the bonding relation is transitive (as co-consciousness seems to be), so that when one microphenomenal state stands in this relation to two other phenomenal states of two other subjects, all three will be jointly experienced by a single subject. If so, then given the ubiquity of spatiotemporal and causal relations, it looks as if the microphenomenal states throughout the universe may stand in this relation, yielding a single giant subject. If on the other hand the relation is not transitive and one has distinct subjects for different instances of the relation, then one will have far too many subjects and it is hard to see how we will get macrosubjects. Perhaps there are intermediate possibilities in which the relation is just nontransitive enough to yield nontrivial macrosubjects, but it is hard to see where this structure will come from.<sup>8</sup>

Perhaps there are intermediate options, but it is not at all easy to see how phenomenal bonding will avoid the Scylla of a universal subject and the Charybdis of fragmentary subjects. To yield human consciousness, we presumably want phenomenal bonding to bond a limited multiplicity

---

<sup>8</sup>Dainton (2011) suggests that a nontransitive view of co-consciousness can help with the combination problem by making it coherent that microsubjects and macrosubjects share experiences, but he does not really address how the relation could be structured to yield a nontrivial structure of macrosubjects.



of microsubjects associated with the human organism, without bonding these to microsubjects elsewhere. It is not at all easy to see what sort of fundamental microphysical relation has this character. Fundamental spatiotemporal and causal relations do not seem to. Perhaps there are derivative causal relations that have this character (a certain sort of informational integration along the lines of Tononi's hypothesis, perhaps?), but these relations are not themselves fundamental. One might suggest that these derivative relations stand to underlying fundamental relations as the bonding relation stands to a more fundamental proto-bonding relation; but now we have a new combination problem concerning how proto-bonding relations can combine to yield a bonding relation.

One might also worry about the quality combination problem. The co-consciousness relation does not seem to help much here: presumably the limited palate of microqualities experienced by microsubjects will also be experienced by macrosubjects, and it is not clear how a rich tapestry of macroqualities will emerge. Perhaps there is another sort of phenomenal bonding relation such that bonded microqualities yield a novel macroquality with a different character, but this relation must go well beyond co-consciousness, and it is not clear how it will work.

It is also far from clear how phenomenal bonding will help with the structure combination problem. Insofar as our underlying phenomenal relation is the categorical basis of spatiotemporal or causal relations, one would expect it to have the same structure as those relations, and one would expect the bonded systems to have structure isomorphic to the corresponding composite spatiotemporal or causal structure. But that is not what we find. So new insights are needed here.

Still, I think that phenomenal bonding is one of the more promising approaches to the combination problem. I have not begun to canvas all the potential phenomenal relations available to a bonding theorist above, and it is not clear that there is a decisive objection to all such theories (the structural mismatch objection is perhaps the best candidate). So I think phenomenal bonding theories are well worth attention.

## **6.4 Deflating awareness**

Another approach focuses on the awareness relation that subjects stand in to qualities. This relation plays a particularly crucial role as it is arguable that all conscious experience consists in a subject's awareness of qualities. As such, if we can explain how microexperiences and microphysics constitute each instance of the awareness relation between subjects and qualities, we will have solved the combination problem.

It is easy to think of the awareness relation as a primitive relation, on which case it is hard to see how instances of it could be constituted by more basic entities. So a constitutive panpsychist or panprotopsychist may need to deny that it is a primitive relation, and explain how instances of it can be constituted.

As with subjects, an extreme deflationary strategy here involves eliminativism: the denial that there is any awareness in experience. This is the strategy famously taken by James in “Does ‘Consciousness’ Exist?”. He suggests that in experience we find only qualities, with no subjects and no relation of awareness. This view certainly makes the combination problem easier to solve. It has not proved popular, however. It seems introspectively obvious that we are aware of qualities (indeed, I think we are aware of our awareness of qualities; see Chalmers 2013 for an argument). Further, our awareness of qualities plays a natural role in explaining our knowledge of qualities. We can conceive of a situation with qualities that no one is aware of, but such a situation seems very different from ours.

A more moderate deflationary strategy is to endorse some sort of reductionism about the awareness relation. One sort of strategy is to give a causal or functional analysis of awareness. For example, perhaps to be aware of a quality is to stand in a certain causal relation to instances of it, or perhaps it is to have states that play a certain functional role associated with that quality. Given this much, awareness (as a relation between organisms and qualities, say) might be grounded in physical terms alone, or in terms of physical states plus qualities. A version of this strategy is taken by Coleman (2012), who uses a functional account of awareness along with instances of qualities to ground awareness of those qualities.

The obvious objection here is that the same considerations that motivate the rejection of functionalism about experience also motivate the rejection of functionalism about the awareness relation. Awareness involves phenomenology, and there are good reasons to think that no mere functional state can constitute phenomenology. For example, one can conceive of any such functional state in the absence of phenomenology, and in particular in the absence of awareness.

Perhaps there are less deflationary accounts of the awareness relation on which it still can be the result of combination. For example, perhaps awareness in microsubjects could somehow constitute awareness in macrosubjects, or protophenomenal properties involving proto-awareness could somehow constitute awareness. It is far from obvious just how this will work, however.

A view like this has the potential to answer the subject combination problem. Anything that is aware of a quality is a subject, so if this approach can show how brains or organisms stand in the awareness relation to qualities, then it will show how brains or organisms can be subjects. On

the other hand, the fact that awareness requires subjects might simply suggest that the awareness combination problem is just as hard as the subject combination problem and is subject to the same sort of worries.

The view does not say much about the quality combination problem: it presupposes qualities rather than explaining them. It has the potential to say something about the structural combination problem, by seeing phenomenal structure as the awareness of complex structured qualities. If awareness of those qualities can be explained, phenomenology will be explained. Still, it is not easy to explain awareness of complex structured qualities starting from a base whose structure is quite different.

## 6.5 Panqualityism

A historically popular form of Russellian monism is what Herbert Feigl (1958) called “panqualityism”. This is a view on which the quiddities associated with microphysical properties are *qualities*. Qualities are not phenomenal properties. Rather, they are perceived qualities: the properties we are aware of in experience, such as qualitative redness, greenness, squareness, and so on. Arguably for every quality Q, there is a phenomenal property consisting in awareness of Q, and vice versa.

As such, panqualityism is a form of panprotopsychism. Because qualities are so closely related to phenomenal properties, however, this form of panprotopsychism is closely related to panpsychism. Like other forms of panprotopsychism, it can also be seen as a sort of neutral monism. Indeed, something like this seems to have been the preferred view of neutral monists such as Mach, James, and Russell. It has recently been revived in this guise by Sam Coleman (2013).

I have discussed panqualityism at length in “Panpsychism and Panprotopsychism”, so I will discuss it only briefly here. Panqualityism (like other forms of panprotopsychism) has the advantage that it has no microsubjects at the basic level, so it avoids James’s subject-summing problem. Still, as the view stands, it seems to leave all three main strands of the combination problem open. It is unclear how microqualities can constitute a macrosubject, or how they can constitute macroqualities, or how they can constitute the structure of macroexperience. One needs one of the other solutions to handle each of these issues.

To handle subjects, the historical neutral monists appealed to deflationism (perhaps eliminativism) about subjects and deflationism (perhaps eliminativism) about awareness. More recently, Coleman has appealed to functionalism about awareness here. I think the objections in the previous sections apply strongly here. For example, one can use a conceivability argument (as I do in

“Panpsychism and Panprotopsychism”) to make the case for an explanatory gap between qualities and awareness, and so between qualities and experience.

In addition, panqualityism does not have much that is distinctive to say about the quality combination problem or the structure combination problem, though perhaps it could adapt elements of other proposals here (and see Coleman, this volume). Overall, it seems to me that while panqualityism is an interesting view, it is not obviously more promising than panpsychism in addressing the combination problem.<sup>9</sup>

## 7 The quality combination problem

So far I have focused mainly on the subject combination problem. I turn now to the quality combination problem. How do microqualities combine to yield macroqualities? And what in particular of the palette problem: the worry that a small palette of fundamental microqualities cannot generate the vast array of macroqualities that we find in experience?

Qualities here need not be understood as perceived qualities, as in the previous section. Qualities in that sense need not be instantiated in experience. What is instantiated are phenomenal qualities, which involve awareness of perceived qualities. So what we really need to explain is how a small palette of microphenomenal (or protophenomenal) properties can generate awareness of a vast array of macrophenomenal properties. It is quite plausible that principles for combining perceived qualities will play a role in explaining principles for combining phenomenal qualities, but the matter remains open. So I will think about combination both for perceived qualities and for phenomenal qualities.

The initial issue here is whether qualities can combine to yield other qualities at all. We understand how this can work when perceived qualities are coinstantiated. If the same object simultaneously instantiated a degree of redness and a degree of whiteness (at the same location), it

---

<sup>9</sup>A sixth idea to address the subject combination problem, proposed by Luke Roelofs (2014), is that of *mereological inheritance*: composite entities inherit experiences from subjects that are their parts. Roelofs proposes this in either a “conditional” version where the subjects must meet certain further conditions (e.g. being appropriately related) or a “basic” version where any composite inherits the experiences of its parts (perhaps because all fundamental properties are inherited by wholes from parts). The obvious problem for most versions is that inheritance principles of either sort do not seem to be a priori: it seems that one can straightforwardly conceive of the relevant microsubjects without any inheritance by macrosubjects. If so, then for this view to yield a version of constitutive panpsychism, the inheritance principles will have to be strong a posteriori necessities. The same worry applies to other elements of Roelofs’ interesting and comprehensive framework for dealing with the combination problem.

will instantiate pinkness. Something similar goes for coinstantiated phenomenal qualities. If the same entity simultaneously is aware of a degree of redness and aware of a degree of whiteness (at the same location), it is plausibly aware of pinkness (at that location). But in general separately instantiated qualities (the redness and whiteness of distinct objects) do not yield a combined quality, and nor do separately instantiated phenomenal qualities. So we need a model of how combination of qualities can work.

This issue may interact with the issue of whether high-level awareness is constituted by low-level awareness. Perhaps such an account can explain how awareness of two distinct qualities by two distinct entities in a complex system can yield awareness of entirely distinct qualities. But it is not at all clear how this will work, especially if we reject highly deflationary accounts of awareness,

What about the palette problem? The two main classes of solutions here are *small-palette* solutions and *large-palette* solutions. Small-palette solutions argue that all macroqualities can be generated from just a few microqualities, if we find the right underlying microqualities with sufficient flexibility and generality. It is far from obvious that such a class can be found, but it is also not obviously out of the question. Small-palette solutions are very much subject to the previous problem of how quality combination works, however.<sup>10</sup>

Large-palette solutions suggest instead that the full range of macroqualities are included among the microqualities. So there are microqualities associated with different colors, sounds, smells, tastes, and so on. A sufficiently rich large-palette solution might eliminate the need for quality combination altogether, thereby removing the problem of how quality combination works, or at least reducing it to the issue of how macrosubjects can inherit (awareness of) qualities from microsubjects.

The cost is that the plethora of qualities raises familiar problems of mental causation. On a Russellian view, microqualities are causally efficacious in virtue of serving as quiddities for microphysical properties. Given that there are only a few fundamental microphysical properties and one quiddity for each of these, there can be only a few microphenomenal quiddities. So only a few microqualities can be causally efficacious, and the rest will be epiphenomenal.

---

<sup>10</sup>Roelofs (2014) outlines a small-palette view of quality combination in terms of operations whereby microqualities are “confused” and “refracted” into macroqualities by high-level cognitive processes (see also Coleman, this volume, for a related view in terms of “contamination”). Prima facie this proposal leaves the usual explanatory gaps (one can conceive of the low-level qualities and the physical dynamics without any refraction into high-level qualities), so again it is not easy to see how it works as a sort of constitutive panpsychism without an appeal to strong a posteriori necessities.

A large-palette proponent might suggest that microphysical properties can be multiply realized by many different quiddities, but this greatly complicates the simplicity of the standard Russellian view.<sup>11</sup> The suggestion requires that the apparent simplicity of physics (with a small number of fundamental properties and laws) is in fact concealing a much more complex underlying level with a vast multiplicity of fundamental properties and fundamental laws, all connected in such a way to yield the appearance of simplicity. Alternatively the large-palette proponent might reject the Russellian view and deny that microqualities are quiddities, but then they will need another way to make the microqualities causally efficacious. If they allow the microqualities to interfere with microphysical dynamics, this will tend to lead back to a quasi-Russellian view with a much more complicated dynamics. Large-palette solutions seem once again to be stuck with either a form of epiphenomenalism or radical revisions to the fundamental dynamics of the physical world.

## 8 The structure combination problem

What about the structure combination problem: how can microphenomenal and microphysical structure yield macrophenomenal structure? Recall that the structural mismatch argument was presented earlier as an apparently inconsistent tetrad of claims. With a little elaboration we can turn this tetrad into a direct argument against constitutive Russellian panpsychism.

- (1) If Russellian panpsychism is true, microphenomenal structure is isomorphic to microphysical structure.
- (2) If constitutive panpsychism is true, microphenomenal (and microphysical) structure constitutes macrophenomenal structure.
- (3) Microphysical structure constitutes only macrophysical structure.
- (4) If microphenomenal structure is isomorphic to microphysical structure, then any structure constituted by microphenomenal structure (and microphysical structure) is isomorphic to a structure constituted by microphysical structure.
- (5) Macrophenomenal structure is not isomorphic to macrophysical structure.

---

(6) Constitutive Russellian panpsychism is false.

---

<sup>11</sup>Tom McClelland has suggested a multiple-realization version of a large-palette view to me, while Pat Lewtas has suggested a non-Russellian version.

Here structure is understood as quasi-mathematical structure involving both internal complexity of states and the quality spaces that they fall into.

Premise (1) is a consequence of the thesis that the quiddity associated with a microphysical property is isomorphic to that property. For example, if mass has a scalar structure, the associated quiddity (what plays the mass role) has a scalar structure. If charge has a binary structure, the associated quiddity (what plays the charge role) has a binary structure. On many Russellian views, microphysical properties such as mass and charge are identical to the associated phenomenal (or protophenomenal) quiddity, in which case they are guaranteed to have the same structure. But even if the two are distinct, one can still expect that in order to be able to play the mass role, a quiddity must have the scalar structure associated with mass.

Premise (2) is close to being true by definition, and premise (3) is highly plausible (perhaps even also true by definition). Premise (4) is in the face of it a plausible general principle about structure. Premise (5) is also highly plausible: the macrophenomenal structure of my visual field is *prima facie* very different from the macrophysical structure of my brain, and it will often (for example in cases of illusion) be quite different from the macrophysical structure of other parts of the world.

The argument is not irresistible, of course. Premise (1) might be denied by someone who holds that quiddities can have surplus structure over and above that of the associated microphysical properties. For example, where mass has a simple scalar structure, perhaps an associated phenomenal quiddity might involve awareness of a certain degree of redness, which has a more complex relational structure due to the role of awareness. One could also say in the reverse direction that microphysical properties have surplus structure that microphenomenal properties do not, perhaps because some but not all microphysical properties have phenomenal quiddities. Still, given a Russellian view, it is not easy to see how these structures could be so different that they yield the vast differences between macrophysical and macrophenomenal structure.

Premise (2) could be denied by someone who says that macrophenomenal structure is constituted by microphenomenal (and microphysical) *qualities*, where these qualities go beyond microphenomenal structure. For this view to help with the problem, specific microphenomenal qualities (phenomenal greenness, say) will have to make a difference to the resulting macrophenomenal structure, so that the latter does not straightforwardly correspond to microphenomenal structure alone. It is not easy to see how this nonstructural factor at the micro level could make a structural difference at the macro level, however.

Premise (3) is true by definition on one reading, where the macrophysical is understood as whatever is constituted by the microphysical. One could deny the premise by understanding the macrophysical more narrowly, however, as I will discuss shortly.

Premise (4) appears to be a plausible principle about structure. It might be denied by someone who holds that although microphenomenal and microphysical structure are isomorphic, the rules of composition that apply to the former differ from the rules of composition that apply to the latter. It is not easy to see how this works, however. If microphenomenal and microphysical properties are identical (because mass is identical to the phenomenal property that plays the mass role), it is especially hard to see how a single set of properties could be subject to distinct modes of composition. Even if they are merely isomorphic (because mass is isomorphic to the phenomenal property that plays the mass role), it is hard to see how the two could compose so differently. I return to this issue shortly.

Finally, one could deny premise (5), holding that macrophenomenal structure mirrors macrophysical structure. One route here holds that we are mischaracterizing macrophenomenal structure. Stoljar (2001) suggests that the apparent structure of the visual field is not part of the structure of an experience, but only part of the structure represented by the experience. Still, it is plausible that an experience's representational content is itself part of its structure. Even on a representational view, it is plausible that experiences can be similar or different to each other in a manner isomorphic to the way that their representational contents are similar or different to each other: an experience of red31 is similar to an experience of red32 but dissimilar to an experience of green31, just as red31 is similar to red32 but dissimilar to green31. So the relevant structure seems at least to be an aspect of macrophenomenal structure. If so, premise (5) remains plausible.

Another way to deny premise (5) is to hold that there exist macrophysical structures that are isomorphic to apparent macrophenomenal structures: spatial and qualitative replicas of the visual field, for example. These replicas might exist somewhere in the brain, as a physical-sense-datum theorist or a topographic map theorist might hold, or they might exist in the external world, as a naive realist might hold. Still, the macrophysical structure of topographic maps is sufficiently far from that of the visual field to cause problems for the first view, and cases of illusion and hallucination cause obvious problems for the second view.

Perhaps the best way to respond to the argument is to say that it equivocates on "macrophysical structure". We might say that *narrowly macrophysical* structure is macroscopic structure characterized in terms of physics: for example, in terms of space, time, mass, charge, and so on. *Broadly macrophysical* structure is any structure constituted by microphysics: for example, chemical, bio-



logical, and computational structure. Then a panpsychist can say that premise (3) is true only of broadly macrophysical structure. It is true by definition that microphysical structure constitutes only broadly macrophysical structure, but it is not true that it constitutes only narrowly macrophysical structure, as it constitutes structures that are broadly but not narrowly microphysical. On the other hand, premise (2) is true only of narrowly macrophysical structure: the structure of consciousness is not isomorphic to the spatiotemporal and other narrowly macrophysical structure of the brain, but it may well be isomorphic to other sorts of macrophysical structure there.

Most obviously, one can suggest macrophenomenal structure is isomorphic to certain *information* structure in the brain. For example, the structure of the visual field corresponds to a structure of visual information represented in the brain. I took a version of this line in *The Conscious Mind*. I think something like this has to be the best option for the panpsychist: it seems clear that the structure of the visual field corresponds to information structure in the brain and not to spatial or qualitative structure. The question is whether this line can be made to work.

It is not easy to see how this line can work for a constitutive Russellian panpsychist. From the perspective of physics, high-level information structures are derivative aspects of a more encompassing and more basic narrowly macrophysical structure. We might expect that on a constitutive Russellian view, macrophenomenal properties would have this more basic structure rather than the somewhat arbitrary informational structure. One can bring this out as follows.

On a Russellian view of physics, it is natural to hold that just as there are microquiddities associated with microphysical properties (such as mass), there are macroquiddities associated with narrowly macrophysical properties (such as macroscopic mass). It is also natural to hold that when certain microphysical properties constitute a macrophysical property, the corresponding microquiddities constitutes the corresponding macroquiddity. For example, when microphysical mass constitutes macrophysical mass, the microquiddity of the former constitutes the macroquiddity of the latter. Because the macroquiddity corresponds so closely to the macrophysical property, we should expect them to have isomorphic structure for reasons discussed under premise (1). On a constitutive Russellian view, it is natural to hold that these macroquiddities are macrophenomenal properties, which will then be isomorphic to narrowly macrophysical properties.

At this point, the constitutive Russellian panpsychist may say there are both *narrow macroquiddities*, the quiddities of narrowly macrophysical properties, and *broad macroquiddities*, the quiddities of broadly macrophysical properties, with different macrophenomenal properties playing both roles. Then the macrophenomenal properties we experience might be broad macroquiddities: quiddities of informational properties, for example. This view naturally goes with the view

that while microphenomenal and narrow macrophysical properties are highly natural and play a special role in physics, the macrophenomenal properties we experience (like broadly macrophysical properties) are less natural and more arbitrary from the point of view of physics. Still, it remains unclear just why phenomenal microquiddities should give rise to broad phenomenal macroquiddities, and how these broad phenomenal macroquiddities relate to narrow phenomenal macroquiddities: in particular how can these macroquiddities stand in the constitutive relation that is plausibly required to avoid causal exclusion worries?

I suggested in *The Conscious Mind* that the principles of phenomenal composition might more closely reflect the constitution of information than the constitution of standard macrophysical structure. Again, I think that something like this is perhaps the only viable line for a panpsychist or panprotopsychist. But it is not at all clear why or how phenomenal composition could work this way while still being a sort of constitutive composition. Certainly one could articulate laws of informational structure for phenomenology, but it is not easy to see how these will be metaphysically necessary rather than brute nomic principles.

In any case, if we are looking to either solve the combination problem or to prove it unsolvable, I think the structural mismatch problem is a promising place to focus. It may be that reasoning along the lines I have given here can be made more rigorous to exclude all possible solutions; and it may be that tightening up the reasoning will reveal the avenues that a panpsychist or panprotopsychist may exploit. In any case, it is clear that the structural mismatch argument is a significant challenge that all Russellian monists must answer.

## 9 Conclusion

What, then, are the prospects for solving the combination problem? On my view, the avenues that seem to be perhaps the most worth exploring are phenomenal bonding or quantum holism (to solve the subject combination problem), small qualitative palettes (to address the quality combination problem), principles of informational composition (to address the structure combination problem), and a somewhat deflationary account of awareness of qualities to tie all these aspects together. It is not at all clear whether these ideas can work together in such a way that all of the combination problems are solved at once, however.

After a close analysis of the many aspects of the combination problem and the limited resources for solving them, it is easy to be pessimistic about the prospects for a solution. What emerges is that panpsychism and panprotopsychism, at least in their constitutive Russellian form,

are subject to extraordinary constraints in finding a theory of consciousness. It is hard enough to find a theory of consciousness that works on dualist terms, where we are allowed to take macro-subjects and macrophenomenal properties as primitive and appeal to numerous contingent psychophysical laws. The Russellian monist is constrained to find a theory whereby macroexperience is constituted by a tiny range of underlying primitive properties and without any further contingent fundamental laws. This is a little like trying to juggle seven balls in the air with both hands tied behind one's back.

It may be that the constraints imposed by the combination problem are so strong that the challenge cannot be answered. Or it may just be that trying to satisfy the constraints will point someone toward the correct form for a fundamental theory of consciousness.

## References

Chalmers, D.J. 2011. Frege's puzzle and the objects of credence. *Mind* 120:587-635.

Basile, P. 2010. It must be true, but how can it be? Some remarks on panpsychism and mental composition. *Royal Institute of Philosophy Supplement* 85 (67):93-112.

Blamauer, M. 2011. Taking the hard problem of consciousness seriously: Dualism, panpsychism and the origin of the combination problem. In (M. Blamauer, ed.) *The Mental as Fundamental*. Ontos Publishing House.

Chalmers, D.J. 2003. Consciousness and its place in nature.

Chalmers, D.J. 2009a. Ontological anti-realism. In (D.J. Chalmers, D. Manley, and R. Wasserman, eds) *Metametaphysics: New Essays in the Foundations of Ontology*. Oxford University Press.

Chalmers, D.J. 2009b. The two-dimensional argument against materialism.

Chalmers, D.J. 2012. *Constructing the World*. Oxford University Press.

Chalmers, D.J. 2013. The contents of consciousness: Reply to Hellie, Peacocke, and Siegel. *Analysis*.

Coleman, S. 2012. Mental chemistry: Combination for panpsychists. *Dialectica* 66:137-66.

Coleman, S. 2013. The real combination problem: Panpsychism, micro-subjects, and emergence. *Erkenntnis*.

Coleman, S. this volume. Panpsychism and neutral monism: How to make up one's mind. In (L. Jaskolla and G. Bruntrup, eds) *Panpsychism*. Oxford University Press.

Dainton, B. 2011. Review of *Consciousness and its Place in Nature*. *Philosophy and Phenomenological Research* 83:238-261.

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'*. University of Minnesota Press.

Gabora, L. 2002. Amplifying phenomenal information: Toward a fundamental theory of consciousness. *Journal of Consciousness Studies* 9 (8):3-29.

Goff, P. 2006. Experiences don't sum. *Journal of Consciousness Studies* 13 (10-11):53-61.

Goff, P. 2009. Why panpsychism doesn't help to explain consciousness.

Goff, P. forthcoming. *Consciousness and Fundamental Reality*.

Goff, P. 2009a. Can the panpsychist get round the combination problem? In (D. Skrbina, ed) *Mind that Abides: Panpsychism in the New Millennium*. John Benjamins.

Goff, P. 2011. There is no combination problem. In (M. Blamauer, ed.) *The Mental as Fundamental*. Ontos Publishing House.

Goff, P. this volume. The phenomenal bonding solution to the combination problem. In (L. Jaskolla and G. Bruntrup, eds) *Panpsychism*. Oxford University Press.

Hunt, T. 2011. Kicking the psychophysical laws into gear: A new approach to the combination problem. *Journal of Consciousness Studies* 18(11-12): 96-134.

James, W. 1904. Does "consciousness" exist? *Journal of Philosophy, Psychology, and Scientific Methods* 1(18):477-491.

James, W. 1895. *The Principles of Psychology*. Henry Holt.

Lewtas, P. forthcoming. Building minds: Solving the combination problem. Unpublished manuscript.

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

Lockwood, M. 1993. The grain problem. In (H. Robinson, ed) *Objections to Physicalism*. Oxford University Press.

Maxwell, G. 1979. Rigid designators and mind-brain identity. *Minnesota Studies in the Philosophy of Science* 9.

Montero, B.G. this volume. What combination problem? In (L. Jaskolla and G. Bruntrup, eds) *Panpsychism*. Oxford University Press.

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

Roelofs, L. 2014. *Combining Minds: A Defence of the Possibility of Experiential Combination*. Ph.D. thesis, University of Toronto.

Rosenberg, G.H. 1998. The boundary problem for phenomenal individuals. In (Stuart R. Hameroff, Alfred W. Kaszniak & A. C. Scott, eds.), *Toward a Science of Consciousness: The Second Tucson Discussions and Debates*.

Rosenberg, G.H. 2004. *A Place for Consciousness*. Oxford University Press.

Rosenberg, G.H. 2014. Causality and the combination problem. In (T. Alter and Y. Nagasawa, eds.) *Consciousness in the Physical World: Perspectives on Russellian Monism*. Oxford University Press.

Russell, M. 1921. *The Analysis of Mind*.

Russell, B. 1927. *The Analysis of Matter*.

Seager, W.E. 1995. Consciousness, information, and panpsychism *Journal of Consciousness Studies* 2:272-88.

Seager, W.E. 2010. Panpsychism, aggregation and combinatorial infusion. *Mind and Matter* 8:167-84.

Seager, W.E. this volume. Panpsychist infusion. In (L. Jaskolla and G. Bruntrup, eds) *Panpsychism*. Oxford University Press.

Shani, I. 2010. Mind stuffed with red herrings: Why William James' critique of the mind-stuff theory does not substantiate a combination problem for panpsychism. *Acta Analytica* 25 (4):413-434.

Skrbina, D. 2011. Mind space: Toward a solution to the combination problem. In (M. Blaumauer, ed.) *The Mental as Fundamental*. Ontos Publishing House.

Stoljar, D. 2001. Two conceptions of the physical. *Philosophy and Phenomenological Research* 62:253-81.

Strawson, G. 2006a. Realistic monism: Why physicalism entails panpsychism. *Journal of Consciousness Studies*.

Strawson, G. 2006b. Panpsychism? Reply to Commentators with a Celebration of Descartes. *Journal of Consciousness Studies* 13: 184-280.

Tononi, G. 2008. Consciousness as integrated information: A provisional manifesto. *The Biological Bulletin* 215: 216-242.