

# Panpsychism's combination problem is a problem for everyone\*

Angela Mendelovici

amendel5@uwo.ca

Penultimate draft

## Abstract

The most pressing worry for panpsychism is arguably the *combination problem*, the problem of intelligibly explaining how the experiences of microphysical entities combine to form the experiences of macrophysical entities such as ourselves. This chapter argues that the combination problem is similar in kind to other problems of mental combination that are problems for everyone: the problem of phenomenal unity, the problem of mental structure, and the problem of new quality spaces. The ubiquity of combination problems suggests the *ignorance hypothesis*, the hypothesis that we are ignorant of certain key facts about mental combination, which allows the panpsychist to avoid certain objections based on the combination problem.

## 1 Introduction

*Panpsychism* is the view that the phenomenal experiences of macrophysical items, like ourselves, are nothing over and above combinations of phenomenal

---

\*Mendelovici, Angela. (2019). Panpsychism's Combination Problem Is a Problem for Everyone. In William Seager (ed.), *The Routledge Handbook of Panpsychism*. London, UK: Routledge.

experiences of microphysical items, where the relevant modes of combination might include physical properties and relations.<sup>1</sup> Most versions of the view can be seen as being motivated by the perceived failure of *physicalism*—the view that consciousness is nothing over and above some arrangement of (non-experiential) physical items—to provide an intelligible explanation of phenomenal consciousness, together with a desire to explain at least our own experiences in more fundamental terms. Physicalist attempts at explaining consciousness in terms of fundamental non-experiential physical reality are subject to explanatory gap worries (Levine 1983), the conceivability argument (Chalmers 1996), and the knowledge argument (Jackson 1982), all of which arguably arise from physicalism’s failure to render intelligible the putative connection between phenomenal consciousness and physical reality. *Dualism*, which takes phenomenal experiences such as our own to be fundamental, avoids such worries by denying that phenomenal experiences can be explained in terms of something else, but gives up on the reductive spirit of physicalism, taking our phenomenal experiences to be primitive, and perhaps brute and inexplicable, features of reality.

Panpsychism attempts to get the best of both worlds, combining physicalism’s reductive spirit with dualism’s skepticism about explaining consciousness in non-experiential terms. Like physicalism, panpsychism aims to explain our phenomenal experiences in terms of something else, though it denies that this something else is wholly non-experiential. Like dualism, panpsychism takes at least some instances of phenomenal consciousness to be fundamental. *Our* experiences may not be fundamental, but they are made up of experiences that are.

Unfortunately, it is not clear that panpsychism can offer an intelligible explanation of the phenomenal experiences of macrophysical entities like ourselves at all, and so it is not clear that panpsychism is any better off than physicalism with respect to explaining our experiences. The problem is that it is not clear

---

<sup>1</sup>This is what Chalmers (2016) calls “constitutive panpsychism”. “Panpsychism” is sometimes more generally defined as the view that consciousness is fundamental and ubiquitous.

how fundamental experiences can come together to form experiences such as our own. This problem is the combination problem, and it has been discussed at length by Seager (1995), Goff (2006), Stoljar (2006), Basile (2010), Coleman (2012), Roelofs (2014), Chalmers (2016), Morch (2014), and others.

The aim of this chapter is to clarify the combination problem, assess the extent to which problems of mental combination are unique to panpsychism, and consider the implications for arguments against panpsychism. I will argue that the panpsychist's combination problem might not be hers alone and that this suggests an "epistemic" reply to objections to panpsychism from the combination problem.

## 2 Panpsychism and the combination problem

Panpsychism is a theory of *phenomenal consciousness*, the felt, qualitative, subjective, or "what it's like" (Nagel 1974) aspect of mental life. We can call particular instances of phenomenal consciousness (*phenomenal experiences*), and the specific "what it's like" or felt quality of an experience its *phenomenal character*. For example, an experience of redness might be said to have a "reddish" phenomenal character.

According to panpsychism, the fundamental physical constituents of reality (*microphysical* entities) have experiences, and the experiences of non-fundamental physical items (*macrophysical* entities) are constituted by the experiences of microphysical items, perhaps combined in a certain way, where the relevant mode of combination might involve functional and physical properties and relations. We can call the experiences of microphysical items *microexperiences* and the experiences of macrophysical items *macroexperiences*. For panpsychism, phenomenal consciousness is both a posit and an explanandum: panpsychism aims to explain macroexperiences such as our own and it does so by positing microexperiences.

Given that a central motivation for panpsychism is the failure of physicalism

to provide an intelligible explanation of phenomenal consciousness, I will assume that panpsychists aim to provide an explanation of macroexperiences that is *intelligible*. I will take this to require that the macroexperiential facts are a priori entailed by the facts about microexperiences and how they are combined. I will not assume, however, that panpsychism requires that we can ever *know* such a theory, and I will eventually suggest that such a theory might not be knowable by us.

Perhaps the most pressing worry for panpsychism is the *combination problem*, the problem of explaining how the hypothesized microexperiences combine to form macroexperiences, such as our own observed experiences. We can sharpen the worry with some assumptions:

- (A1) Macroexperiences are not identical with any one of their constituent microexperiences.
- (A2) Macroexperiences are had by subjects that are distinct from the subjects of any one of their constituent microexperiences.
- (A3) Macroexperiences have phenomenal characters that are not had by any of their constituent microexperiences.

Given these three assumptions, the combination problem becomes that of explaining how groups of microexperiences come together to constitute (1) *new experiences*, which belong to (2) *new subjects*, and have (3) *new phenomenal characters*. We can thus tease apart three combination problems for panpsychism:

- (CP1) The new experience problem
- (CP2) The new subject problem
- (CP3) The new phenomenal characters problem

Note that, given our definition of panpsychism, none of the assumptions that give rise to the combination problems form a definitional part of panpsychism,

and so a panpsychist solution to these problems might coherently deny any one of them.

Problems (CP1) and (CP2) are sometimes lumped together under the heading of “the subject combination problem” and taken to be the central or most difficult part of the combination problem (see Roelofs’ contribution to this volume). As we will soon see, (CP1) and (CP2) are distinct problems, though they interact with one another in interesting ways.<sup>2</sup>

The remainder of this section elaborates upon the combination problems for panpsychism and suggests that what makes them particularly challenging is that they require mental things to come together to form more than a mere collection of their parts.

**The new experience problem.** The *new experience problem* is the problem of explaining how microexperiences combine to form distinct macroexperiences. For example, according to panpsychism, two microexperiences, e1 and e2, when combined in the right way, might give rise to a distinct macroexperience, E. The problem is that of explaining how this new experience arises. What makes the new experience problem challenging is that it is not clearly intelligible why a collection of experiences, however organized, should result in a further experience.

The new experience problem can be avoided by rejecting assumption (A1), the assumption that microexperiences combine to form distinct macroexperiences, and instead claiming that each macroexperience is identical to a constituent microexperience. On such a view, macroexperiences are present at the fundamental level, and so there are no “new” experiences to account for. Leibniz’s (1714/1989) monadology is such a version of panpsychism. One worry with this general approach is that it seems there would be a surprising structural

---

<sup>2</sup>Chalmers (2016) distinguishes between three combination problems: the subject combination problem, the quality combination problem, and the structure combination problem. (CP1) and (CP2) correspond to Chalmers’ subject combination problem and (CP3) roughly corresponds to Chalmers’ quality and structure combination problems.

Problems like the grain problem (see, e.g., Maxwell 1979, Lockwood 1993, and Stoljar 2006) are not combination problems, as I am understanding them, since they are not problems of explaining how new items combine, but rather worries with certain solutions to such problems.

mismatch between the microphysical properties of the dominant monad and its corresponding experience (see Chalmers 2016). Another reason to at least dislike such a view is simply that taking our own experiences to be fundamental foregoes much of the explanatory appeal of panpsychism over ordinary dualism, which is that it promises to offer an explanation of our own experiences in terms of something else. For these reasons, the panpsychist probably should not try to avoid the new experience problem by rejecting (A1).

**The new subject problem.** The new subject problem is the problem of explaining how subjects of microexperience combine to form distinct subjects of macroexperience. Suppose  $s_1$  and  $s_2$  are the subjects of experiences  $e_1$  and  $e_2$ , respectively. On most natural versions of panpsychism, when  $e_1$  and  $e_2$  combine to form the new experience  $E$ , this experience is an experience of a new subject,  $S$ , which is distinct from  $s_1$  and  $s_2$ . The new subject problem is that of explaining how  $S$  arises from a combination of  $s_1$  and  $s_2$ . The problem is challenging because it is not clearly intelligible why a mere collection of subjects, however organized, should yield a new subject (see, e.g. Goff 2006, 2009).

The new subject problem can be avoided by rejecting (A2), the assumption that the subjects of macroexperiences are distinct from the subjects of any one of their constituent microexperiences, and instead claiming that the subjects of macroexperiences are simply the subjects of one or more of the constituent microexperiences. In the preceding example, we could say that  $E$  is an experience of  $s_1$ ,  $s_2$ , or both  $s_1$  and  $s_2$ , taken severally. Of these options, the first two seem arbitrary (why should  $E$  be an experience of  $s_1$  rather than  $s_2$ ?), which leaves us with the last option:  $s_1$  experiences  $E$ , and  $s_2$  also experiences  $E$ . But such a view, on which, presumably, every macroexperience is had by all the subjects of all its constituent microexperiences, seems a bit excessive. It also faces the same structural mismatch problem as the Leibnizian view discussed previously. For these reasons, the panpsychist probably should not try to avoid the new subject problem by rejecting (A2).

In the literature on the combination problem, the new experience problem is often lumped together with the new subject problem under the label “the subject combination problem,” which is taken to be the problem of explaining how subjects of microexperiences, with their various experiences, combine to form distinct subjects of macroexperiences with distinct experiences (Seager 1995, Goff 2006, 2009, Chalmers 2016). However, the two problems are independent from one another in that it is possible to have a panpsychist view that faces one problem but not the other. As we saw above, it is a theoretical possibility that when  $s_1$  and  $s_2$  combine, a new experience  $E$  arises, but it is an experience of  $s_1$ ,  $s_2$ , or both  $s_1$  and  $s_2$ , severally, rather than an experience of a new subject  $S$ . On such a view, there is a new experience without a new subject. It is also a theoretical possibility that when  $s_1$  and  $s_2$  combine, a new subject  $S$  is formed, but  $S$ 's experience is numerically identical to the experience of  $s_1$  or  $s_2$ , so no new experience arises. On this view, there is a new subject without a new experience.

**The new phenomenal character problem.** The *new phenomenal character problem* is the problem of explaining how the phenomenal characters of microexperiences combine to form the phenomenal characters of macroexperiences. The problem arises from (A3), according to which macroexperiences have phenomenal characters that their constituent microexperiences do not have. For example, we experience colors, shapes, and feelings of déjà vu, but microphysical items presumably do not have all these kinds of experiences.

We can distinguish between two types of new phenomenal characters that the panpsychist might want to accommodate: *complex* phenomenal characters, which are phenomenal characters that have parts that are also phenomenal characters, and *simple* phenomenal characters, which are phenomenal characters that are not complex. For example, the phenomenal character of an experience of a red square might be complex in that it involves as parts both reddish and squarish phenomenal characters, but the phenomenal character of an experience

of redness might be simple, not involving other phenomenal characters as parts.<sup>3</sup>

The panpsychist faces challenges in accommodating both simple and complex new phenomenal characters. Suppose a macroexperience E has a complex reddish-squarish phenomenal character. According to panpsychism, E's complex phenomenal character is a result of the phenomenal characters of its constituent experiences. Perhaps E is a combination of two experiences, e1 and e2, where e1 has a reddish phenomenal character and e2 has a squarish phenomenal character. The problem is that it is not clear why E should have a reddish-squarish phenomenal character, rather than a reddish phenomenal character alongside a squarish phenomenal character. In other words, it is not clear why e1 and e2's phenomenal characters should combine in E to yield a complex whole, a reddish squarish phenomenal character, rather than simply co-exist as two unrelated simple (or simpler) phenomenal characters, reddishness and squarishness. It is even less clear how new simple phenomenal characters should arise from the phenomenal characters of microexperiences, since they do not even have constituent parts that are also phenomenal characters. There aren't even any candidate phenomenal characters to be combined, let alone a way of intelligibly combining them into a new whole.

The problems can be avoided by rejecting assumption (A3): If microphysical items do have the full range of experiences found in macrophysical items, then there need be no combined phenomenal characters. But it is implausible that the full range of experiences found at the macrolevel is found at the microlevel. Many of the phenomenal characters of macroexperiences appear to be too sophisticated to be found at the microlevel, such as feelings of jealousy or

---

<sup>3</sup>Note that simple phenomenal characters might nonetheless be *externally structured* in that they can have properties pertaining to how they are phenomenally, which might make them similar to or different from other simple phenomenal characters. For example, the phenomenal character *reddishness* might be simple but it might nonetheless have a certain value on the dimensions of hue, saturation, and brightness that make it similar to and different from the phenomenal characters of other color experiences. Since color phenomenal characters are characterized by their values on the dimensions of hue, saturation, and brightness, they can be perspicuously modeled as positions in a *quality space* whose axes stand for these dimensions on which they have a value. See Chalmers 2016 for the distinction between internal and external structure.



cognitive experiences of suddenly grasping a difficult concept. Additionally, and perhaps more persuasively, it is implausible that there are enough kinds of microexperiences to correspond to all the kinds of macroexperiences we can have. Many versions of panpsychism are committed to the view that microexperiences are the inner categorical natures of microphysical items, a view sometimes called *Russellian panpsychism*. Presumably, all tokens of a particular type of microphysical entity have the same inner nature, and hence the same kind of microexperience. But, presumably, there is a limited stock of microphysical items, and, presumably, the number of such items is smaller than the number of types of macroexperiences that macrosubjects have. But then there are simply not enough types of microexperiences to correspond to all the types of macroexperiences. Some types of macroexperiences, then, must be combinations of these limited types of microexperiences.<sup>4</sup>

### 3 Combination problems for everyone

Panpsychism's combination problems are challenging (see especially Goff 2006, 2009, Chalmers 2016), but the panpsychist does not face them alone. They are of the same kind as the problems of explaining phenomenal unity, mental structure, and changes in quality spaces, which are problems for anyone holding certain plausible assumptions.<sup>5</sup>

#### 3.1 The new experience problem is not special to panpsychism

Panpsychism's new experience problem is the problem of explaining how microexperiences come together to form distinct macroexperiences. This subsection

---

<sup>4</sup>See Chalmers 2016 and Roelofs 2014 for elaborations of this reason for taking macroexperience to involve phenomenal characters not found in microexperience.

<sup>5</sup>See also Roelofs 2015, which argues that any physicalist or panpsychic view of consciousness faces problems of mental combination if it allows for two conscious things to be part of a larger system that is conscious and whose consciousness is entirely dependent on theirs. He suggests that the problem could be largely avoided by taking subjects to be simple.

argues that the new experience problem is the same in kind as two other well-known problems, the problem of explaining phenomenal unity and the problem of explaining mental structure. The phenomena of phenomenal unity and mental structure arguably involve experiences coming together to form new experiences in much the same way that panpsychism requires microexperiences to come together to form new macroexperiences.

**The problem of phenomenal unity.** You might now be enjoying various visual, auditory, and cognitive experiences. These experiences are in some sense experienced *together*. In contrast, your experiences and the experiences of other people are not experienced together. *Phenomenal unity* is the phenomenon of experiences being experienced together that is present in the former kinds of cases and absent in the latter kinds.

The *problem of phenomenal unity* is the problem of explaining how and why some experiences are phenomenally unified while others are not. Solving this problem is particularly difficult because it seems that what is required for a group of experiences to be phenomenally unified is something more than their co-occurrence. Something like this is assumed by two influential characterizations of phenomenal unity.

On Bayne and Chalmers' (Bayne 2012 and Bayne and Chalmers 2003) characterization, experiences are phenomenally unified when they are subsumed by a single conscious state, that is, when there is a further experience that includes them both. On this characterization, phenomenal unity involves a *new* experience, one that subsumes the unified experiences.

Similarly, Dainton (2000) characterizes phenomenal unity in terms of co-consciousness, where co-consciousness is not merely a matter of experiences occurring at the same time or place, or even in the same subject, but rather "consists in a relationship between experiences that is itself experienced." (p. 4)<sup>6</sup> On this characterization, the phenomenal unity of e1 and e2 involves an experi-

---

<sup>6</sup>Bayne (2012) suggests that his and Chalmers' conception is compatible with Dainton's.

enced relation between e1 and e2, and the experience of this relation is a new experience, distinct from e1 and e2.

**The problem of mental structure.** Our mental states do not form an undifferentiated whole, or a set of totally isolated distinct mental states, but are instead related and structured in various ways. For example, a visual experience of a red circle does not only involve an experience of redness and an experience of a circle, but also involves these experiences being related in a certain way: The experienced redness qualifies the experienced circle. The *problem of mental structure* is that of explaining how mental states come to be structured in this and other ways.<sup>7</sup>

One instance of the problem of mental structure is the binding problem, which roughly has to do with how the operations of different parts of the brain combine to form a unified experience of a particular object having multiple features. Several “binding problems” have been distinguished (see Revonsuo 1999 and Smythies 1994), but the problem that concerns us here is what we might call the *experience binding problem*, the problem of explaining how distinct experiences that are subserved by distinct neural areas are experienced as pertaining to the same consciously represented object.<sup>8</sup> For example, an experience of a red square might involve an experience of redness and an experience of squareness that are subserved by different neural areas. The experience binding problem is that of explaining how these experiences come together to qualify the same represented object.

Another instance of the problem of mental structure concerns intentional structure. *Intentional contents*, what mental states (or other items) “say,” are directed at, or represent, can be structured in various ways. For example, the propositional content <Lisa loves Sally> might be composed of the contents

---

<sup>7</sup>Note that while mental structure might require phenomenal unity, it is not the same thing as phenomenal unity, since some unified experiences might bear no structural relations to one another, such as a visual experience of a red square and a cognitive experience of thinking that  $2+2=4$ , and different phenomenally unified and mental structured experiences might be structured in different ways.

<sup>8</sup>This is, roughly, Smythies’ (1994) “BM2.”

<Lisa>, <loves>, and <Sally>. The *problem of intentional structure* is that of explaining how intentional states representing a content's constituent contents come together to form a complex structured intentional state representing a complex structured content, rather than, say, a set or list of isolated contents. For example, how is it that <Lisa>, <loves>, and <Sally> combine to form <Lisa loves Sally>, rather than the mere set of contents {<Lisa>, <loves>, <Sally>}?<sup>9</sup>

Mental structure quite plausibly involves new mental states, mental states involving but distinct from the mental states that compose them. For example, suppose M1 and M2 are bound to the same represented object. Then there is a mental state distinct from M1 and M2, consisting of M1 and M2 together and organized in a certain way, i.e., as bound to the same represented object. For example, a thought that Lisa loves Sally involves not only the representation of the contents <Lisa>, <Sally>, and <loves> but also a distinct state representing <Lisa loves Sally>.<sup>10</sup>

If the preceding claims about the problems of phenomenal unity and mental structure are right, then the problem of explaining how experiences combine to form new experiences may not be special to panpsychism. On the reasonable assumption that certain kinds of holism are not true, which we will consider shortly, phenomenal unity involves experiences coming together to form new unified experiences, and mental structure involves experiences or intentional states coming together to form new complex experiences or intentional states,

---

<sup>9</sup>A special case of the problem of intentional structure is the problem of the unity of the proposition, which is that of explaining how objects and properties can combine to form full-blown propositions, rather than lists or sets of non-propositional contents (see Gaskin 2008).

Given certain assumptions about the relationship between phenomenal consciousness and intentionality, the binding problem is also a special case of the problem of intentional structure. If we take phenomenal characters to be identical to, constituted by, or realized by intentional contents, as do some versions of representationalism (Tye 2000, Dretske 1995, and Bourget and Mendelovici 2014) and the phenomenal intentionality theory (Kriegel 2011, Pitt 2004, Bourget and Mendelovici 2016, and Mendelovici 2018), then, on certain reasonable assumptions, the binding problem is simply that of explaining a special kind of intentional structure.

<sup>10</sup>This is quite plausible on the view that there is a language of thought (Fodor 1975), on which intentional states involve distinct representations coming together to form complex representations whose contents are a logical construction of the contents of their parts.

respectively. Of course, panpsychism requires that *micro*experiences combine to form new experiences, whereas phenomenal unity and mental structure only require macroexperiences to combine to form new experiences. But it is not clear that what is required is different in kind.

One might object that there is a way out of this commitment in the case of the problems of phenomenal unity and mental structure that is not available in the case of the panpsychist's new experience problem, so the problems are different in kind. The way out of the problems of phenomenal unity and mental structure is to reject the assumption that when we experience a phenomenally unified or mentally structured whole, we also experience its parts. A holistic view of this sort (see e.g. James 1890) avoids commitment to new experiences by denying that macroexperiences ever combine in the relevant way. What appear to be separable parts of our experiences are in fact mere aspects of the experiences, having no distinct and independent existence, but instead having an existence that depends on the whole of which they are an aspect. For example, on such a view, an experience of a red square does not involve an experience of redness and an experience of squareness. Instead, it only involves an experience of a red square, and redness and squareness are mere aspects of this experience.

However, the panpsychist might similarly avail herself to a "holistic" solution to the new experience problem: She might say that the ultimate constituents of reality are not "small" things, but rather the world as a whole, which has one single experience (at least at a time) with many aspects corresponding to what we take to be our experiences (see Goff 2017). Alternatively, she might maintain that the ultimate constituents of reality are or include subjects like ourselves. Like the way out of the problems of phenomenal unity and mental structure, this strategy involves denying that the relevant sort of mental combination occurs. Such a view still qualifies as panpsychist on our definition, since it still maintains that macroexperiences are nothing over and above microexperiences combined in a certain way—it's just that every macroexperience is identical to a single

microexperience. Unless there is good reason to think that the problems of phenomenal unity and mental structure are particularly amenable to the holistic strategy while the new experience problem is not, the availability of this strategy in their case does not suggest that the new experience problem is different in kind from the problems of phenomenal unity and mental structure.

Another objection to the claim that the new experience problem is the same in kind as the problems of phenomenal unity and mental structure is that in the case of new experiences arising from phenomenal unity and mental structure, the new experiences are experiences of the *same* subjects that experience the combined experiences, whereas in the case of the panpsychist's new experiences, the new experiences are assumed to be experiences of *new* subjects. This suggests that perhaps the way in which microexperiences combine to form new macroexperiences is different from the way in which macroexperiences combine to form new macroexperiences, which would mean that the panpsychist's new experience problem is indeed special to panpsychism. Before responding to this objection, it is helpful to consider the question of whether the new subject problem is special to panpsychism, to which I now turn.

### **3.2 The new subject problem is not special to panpsychism**

Let us first assume a fairly thin notion of subjects on which subjects are sets of phenomenally unified experiences. On this notion, when mental combination results in a new experience, that experience automatically has a subject. For example, once phenomenal unity results in a new experience subsuming or including all the unified experiences, we thereby automatically have a subject for that experience.

On the thin view of subjects, there is no mystery as to why phenomenally unified experiences have subjects: they have subjects simply because they are phenomenally unified and subjects are phenomenally unified experiences. On the

face of it, it might seem that the panpsychist can solve the new subject problem in the same way: when the experiences of microsubjects are phenomenally unified, a new macrosubject comes to exist and experiences the phenomenally unified experiences. The new subject problem, then, can be solved by adopting a thin view of subjects and solving the new experience problem, which is a problem for everyone.

There is a worry, however, which brings us back to the worry raised at the end of the previous subsection: The way subjects combine to form new subjects according to panpsychism and the way phenomenally unified experiences come to form subjects of experiences in the case of phenomenal unity are importantly disanalogous. In a case of panpsychist subject combination, a new subject, *S*, experiences microexperiences *m1* and *m2* combined (i.e., a macroexperience *M*), but, it is natural to assume, *m1* and *m2* are each *also* experienced by a subject distinct from *S*. In contrast, in a case of phenomenal unity, when experiences *e1* and *e2* are phenomenally unified to form experience *E*, it is natural to assume that there is only a single subject of experience, which experiences *e1* and *e2* together (i.e., *E*). If so, then what's responsible for the arising of new subjects on panpsychism cannot be the same thing as what's responsible for phenomenally unified experiences having subjects. The problem is not so much to do with how the new subject arises but rather with what happens to the "old" subjects once combined. In the case of phenomenal unity, the old subjects cease to exist or are subsumed by the new subject. In the case of panpsychist subject combination, the old subjects continue to exist. When microexperiences *m1* and *m2* combine into *M*, there are three subjects (the subject of *m1*, the subject of *m2*, and the subject of *M*), whereas when experiences *e1* and *e2* are phenomenally unified to form *E*, there is only one subject (the subject of *E*, which is also the subject of *e1* and *e2*).

This worry arises from two assumptions, the first of which is natural on panpsychism and the second of which is natural on any picture of phenomenal

unity:

- (A) When experiences combine to form further experiences, they are experienced both together and in isolation.
- (B) When experiences are phenomenally unified, they are experienced together but not in isolation.

We can avoid the worry described above by rejecting either of these assumptions. Let us first consider (A). On this assumption, when  $m_1$  and  $m_2$  are combined to form  $M$ , there is an experience of  $m_1$  in isolation, an experience of  $m_2$  in isolation, and an experience of  $m_1$  and  $m_2$  combined (i.e.,  $M$ ). On the thin view of subjects, this means that there are three subjects of experience, a subject of  $m_1$ , a subject of  $m_2$ , and a subject of  $M$ . The panpsychist might choose to deny (A) and instead claim that when  $m_1$  and  $m_2$  are combined, they are experienced together but not in isolation.

One view of panpsychist combination, the *combinatorial infusion view* (Seager 2010, 2016, Morch 2014), makes precisely such claims. On this view, when microexperiences combine to yield macroexperiences, they fuse together and cease to exist independently. As Seager (2010) puts it, they are “absorbed” or “superseded” by the macroexperience they come to constitute. On this picture, when microexperiences combine, the result is only one subject of experience that experiences the combined microexperiences.

The combinatorial infusion view, and any other panpsychist view that rejects the first assumption, avoids the worry that the problems of explaining subject unity and phenomenal unity are different in kind because they yield different treatments of the old subjects of experience. Indeed, Seager suggests that the combinatorial infusion view might help solve the problem of phenomenal unity:

It is also possible to avoid the worry described above by rejecting (B), the assumption that phenomenally unified experiences are experienced together but not in isolation. Perhaps, instead, when  $e_1$  and  $e_2$  are phenomenally unified,  $e_1$  and  $e_2$  are experienced both together *and* severally. There is an experience of  $e_1$



together with e2 (E), an experience of e1 in isolation, and an experience of e2 in isolation. This option might seem unlikely, since we have no phenomenological evidence that phenomenally unified experiences are also experienced in isolation. But note that there is also no phenomenological evidence against this possibility: It is entirely compatible with an experience of E that there exist isolated experiences of e1 and e2. On the thin view of subjects, there would then be three subjects of experience: the subject of e1, the subject of e2, and the subject of e1 and e2 together. Indeed, Roelofs (2016) suggests that such a view is true and helpful to panpsychism, helping us make sense of how experiences can be shared between distinct microphysical and macrophysical entities.<sup>11</sup>

In sum, the worry that the subject combination required by panpsychism has a different source than whatever results in phenomenally unified experiences having subjects depends on two assumptions, either of which can be rejected. If we accept the thin view of subjects and reject one of these assumptions, the panpsychist's subject combination is plausibly of the same kind as whatever results in phenomenally unified experiences having subjects. The claim that panpsychism faces a *special* problem of subject combination depends on both assumptions being true.

The rejection of either (A) or (B) also allows us to respond to the worry described at the end of §3.1 that there is an important difference between the new experiences required by panpsychism and those required by phenomenal unity and mental structure. The alleged difference is that in the case of new experiences arising from phenomenal unity and mental structure, the new experiences are experiences of the *same* subjects that experience the combined experiences,

---

<sup>11</sup>The denial of (B) amounts to a denial of Dainton's (2000, p. 246) "exclusivity principle," which states that any experience can have only one subject, and an endorsement of Basile's (2010) "sharing principle," which states that an experience can be shared by two psychical wholes. Roelofs (2016) argues that denying the exclusivity principle allows us to block an argument against the possibility of between-subjects (phenomenal) unity, allowing that two experiences e1 and e2 had by distinct subjects can be phenomenally unified if there is a third subject with an experience subsuming both e1 and e2. He further suggests that the possibility of between-subjects unity is helpful to panpsychism. See also Basile 2010 and Coleman 2013 for arguments in favor of something like the exclusivity principle on the basis of a holistic view of phenomenal unity, and Roelofs 2016 for a reply.

whereas in the case of the panpsychist's new experiences, the new experiences are experiences of *new* subjects. But if we accept a thin view of subjects and reject (A), then in both cases, the combined experience is an experience of a single subject that is distinct from the subject of the experiences that form the experience's parts. And if we accept a thin view of subjects and instead reject (B), then in both cases, the combined experience is an experience of a single subject that is also the subject of the experiences that form the experience's parts. So, the cases are not disanalogous. Again, the worry that panpsychism faces a *special* problem of mental combination concerning new experiences depends on both assumptions being true.

I have argued that the panpsychist faces no special problem in accounting for new thin subjects of macroexperience. But what if we think that there are such things as subjects on a thicker notion of subjecthood, perhaps one on which subjects can persist over time or exist without having experiences? A theory of thick subjects might reduce them to something else, perhaps to sets of thin subjects that meet certain further criteria, or it might take thick subjects to be *sui generis*. If the panpsychist accepts that there are such thick subjects and that they can combine to form new thick subjects, then, depending on what exactly they are supposed to be, she might face special problems in accounting for the required kind of combination. But I want to suggest that even if the panpsychist accepts that macroexperiences have thick subjects, she need not accept that *microexperiences* have thick subjects that combine to form them. It is enough for the panpsychist to say that microexperiences have thin subjects, and that thick subjects, if there are any, arise in some other way at the macrolevel. The problem of explaining how they arise at the macrolevel, of course, is a problem for anyone who accepts them.

### **3.3 The new phenomenal characters problem is not special to panpsychism**

If the preceding arguments are sound, the new experience and new subject problems are not special to panpsychism. Things are less clear in the case of the new phenomenal characters problem. Recall that there are two types of new phenomenal characters that our macroexperiences seem to exhibit that we need to explain: complex and simple phenomenal characters. Let us consider each in turn.

An example of a complex phenomenal character is the phenomenal character of visually experiencing a red square. This phenomenal character is complex in that it has parts that are also phenomenal characters: the phenomenal character of redness and the phenomenal character of squareness. However, it is more than just a collection of the phenomenal characters of redness and squareness. The two phenomenal characters qualify each other: the redness is experienced in a square shape, and the square shape is experienced as red. For those who believe in cognitive phenomenology, similar examples are available in the case of thought: a conscious thought that Lisa loves Sally might involve the phenomenal characters corresponding to the concepts of Lisa, loving, and Sally, but it is more than just a collection of those phenomenal characters. This is evidenced by the fact that it has a different phenomenal character than the conscious thought that Sally loves Lisa.

To explain how macroexperiences can have new complex phenomenal characters we must explain how complex phenomenal characters arise from their simpler parts. If the phenomenal characters of the simplest parts are those of microphysical entities, then that is all we must do. If it is not, then there is the further problem of explaining how these simple parts arise from the phenomenal characters of microphysical entities, which calls for an explanation of how macroexperiences can come to have new simple phenomenal characters, the second type of new phenomenal character the panpsychist should accommodate.

Let us start with the problem of explaining how complex phenomenal characters arise from their simpler parts. This problem is of the same kind as the problem of mental structure, the problem of explaining how phenomenal and intentional mental features come to be structured. Structured experiences and intentional states have complex phenomenal characters and intentional contents, respectively, which presumably are combinations of their constituent phenomenal characters or intentional contents.

Of course, since the panpsychist but not the non-panpsychist requires that there be microexperiences that combine in the relevant ways, she might require that there be more instances of mental structure than the non-panpsychist, and so her problem might be wider in scope. Still, the problems are of the same kind.

The situation is less clear when it comes to accounting for the combination of phenomenal characters into new *simple* phenomenal characters. The problem of explaining simple combined phenomenal characters is arguably the hard nut, and perhaps the special nut, of the combination problem. The problem seems *hard* because what it seems to require, simple yet combined items, seems incoherent. The problem seems *special* to panpsychism since the non-panpsychist appears not to be committed to such simple yet combined phenomenal characters. She might accept that the simple phenomenal characters in question exist but deny that they are the results of combinations of other phenomenal characters.

The panpsychist might attempt to sidestep this problem of accounting for the combination of phenomenal characters into new *simple* phenomenal characters by denying that macroexperiences have simple phenomenal characters. It might appear that they do, but we are mistaken. For example, it might seem that a reddish phenomenal character is a simple phenomenal character, but it is in fact complex. Roelofs (2014) considers such a view, suggesting that our apparently simple phenomenal characters might be blends of the “alien” phenomenal characters of microexperiences.

In defense of this view, Roelofs points to examples of macroexperiences

that appear simple but plausibly are complex blends of other macroexperiences, such as the apparently simple phenomenal characters of color experiences. An orangish phenomenal character might appear simple, but, he claims, it is in fact a blend of a reddish and a yellowish phenomenal character. Chalmers (2016) similarly gives the example of a pinkish phenomenal character being a blend of reddish and whitish phenomenal characters. Roelofs suggests that such examples show that it is possible for phenomenal characters to blend, and, further, that we are bad at recognizing such blends. In the case of color experience, the reason we can come to appreciate the relevant blends is that we can come to have experiences with the constituent phenomenal characters on separate occasions. For example, we can have experiences with reddish phenomenal characters, and by comparing our reddish experiences with our orangish experiences, we can come to appreciate that “there’s a little bit of red in orange.” In the case of the alien phenomenal characters of microexperiences that blend to form the phenomenal characters of macroexperiences, we are not able to experience the alien phenomenal characters in isolation, so we are not in a position to appreciate that the phenomenal characters of our macroexperiences are blends of them.

However, it is not clear that Roelofs’ examples are effective. An orangish phenomenal character is similar to reddish and yellowish phenomenal characters, but the reason for this similarity isn’t that it is *composed* of them. The phenomenal characters of color experiences might be simple but have various properties that are related to those of other phenomenal characters and that account for the similarities between them, namely their values on dimensions of hue, saturation, and brightness. If this is right, then it is not clear that the panpsychist can avoid commitment to new simple phenomenal characters, and the new phenomenal characters problem remains.

I want to suggest that the problem may not be special to panpsychism. There is a nearby problem facing everyone, that of explaining how we can come to have macroexperiences with new simple phenomenal characters that in some sense

“build on” the phenomenal characters of other macroexperiences:

As we develop and learn, we acquire abilities to have new experiences. For example, a budding wine taster might gradually acquire new abilities to have new wine tasting experiences, such as experiences with fruity, oily, and tannin-ish phenomenal characters. The new phenomenal characters we are able to have in such cases are not wholly unrelated to the phenomenal characters we were previously able to have, but, instead, are similar and different to them in certain ways. We can perspicuously model such relationships of similarity and difference between phenomenal characters using *quality spaces*, abstract spaces with one or more dimensions corresponding to the dimensions of possible variation in a system of phenomenal characters, where different phenomenal characters are represented by different positions in the space. For example, since colors vary in hue, saturation, and brightness, a quality space with axes corresponding to hue, saturation, and brightness is a perspicuous way of modeling them and their similarity relations.

We can think of learning and development as building upon or expanding our pre-existing quality spaces. For example, the wine taster’s quality space for wine-related experiences might expand to include new dimensions. In this way, newly acquired abilities to experience new phenomenal characters might be thought to build upon pre-existing abilities. Call the problem of explaining how exactly the quality spaces characterizing our abilities to have experiences change in such ways the *changing quality space problem*.

On the face of it, the panpsychist’s problem of explaining new simple phenomenal characters and the changing quality space problem seem quite alike: they both require explaining how we can come to experience (at least sometimes) simple phenomenal characters that are not present in our other concomitant or past experiences but that are nonetheless in some important way related to them. Perhaps, then, both problems involve the same kind of mental combination.

Against this, one might suggest that only the panpsychist’s problem is a

problem of mental combination. The panpsychist assumes that an experience's new simple phenomenal characters are a matter of the combination of the phenomenal characters of some constituent experiences, but a solution to the changing quality space problem need not make such an assumption. One non-combinatorial solution to the changing quality space problem maintains that it is macroexperiences' functional roles that determine their specific phenomenal characters. Perhaps, for instance, the functional roles of color experiences fix their phenomenal characters, and when we acquire new concepts, their functional roles, including those in relation to old experiences, alter our quality spaces, allowing for new phenomenal characters.

Even if such a functionalist solution to the changing quality space problem can succeed, this is not automatically a problem for the claim that the panpsychist does not face a special problem in accounting for new simple phenomenal characters, since she can co-opt the functionalist's solution. The panpsychist wants to explain new simple experiences in terms of mental combination, but the relevant modes of combination can include functional properties. Where the non-panpsychist might say that macroexperience E has a new simple phenomenal character C in virtue of playing a certain functional role, R, the panpsychist can say that macroexperience E has a new simple phenomenal character C in virtue of being constituted by experiences e1 and e2, which, together, play functional role R. In effect, the panpsychist can turn the functionalist's non-combinatorial solution to the changing quality space problem into a combinatorial solution for the problem of explaining new simple phenomenal characters. In the same way, other non-combinatorial solutions to the changing quality space problem might be co-opted by the panpsychist. (Of course, this takes some of the bite out of panpsychism, but the view still qualifies as a version of panpsychism.)

I am doubtful, however, that the problem of explaining quality space changes like the ones described above can be solved without appeal to phenomenal ingredients. Let us return to the functionalist proposal, which is arguably the

most promising alternative approach. The problem is that functionalism faces well-known indeterminacy worries. For instance, a set of states that implements a symmetrical system of functional roles could equally well be said to realize at least two quality spaces (see Block 1978 and Palmer 1999). More generally, even if functional roles can determine the relations between phenomenal characters, it is far from clear that there is only one set of phenomenal characters whose members can bear those relations to one another.<sup>12</sup>

The functionalist might attempt to avoid indeterminacy worries by taking at least some functional states to be broad, involving relations beyond the experiencing individual, as on some versions of representationalism (see, e.g., Harman 1990 for this strategy), but this would result in externalism about phenomenal consciousness, the view that a subject's experiences are at least partly determined by environmental features, which is arguably implausible.<sup>13,14</sup> A second strategy is to throw phenomenal characters into the mix. If at least some positions in a quality space have their phenomenal characters independently of their functional roles, then they can serve as "anchor points" (Graham et al. 2007, p. 479), helping to constrain the possible phenomenal character assignments to the rest of the space.<sup>15</sup> However, it is not clear that this is enough to solve indeterminacy worries (see Bourget MS).

If, as I've very briefly suggested above, there are no viable non-combinatorial solutions to the changing quality space problem, then it might just turn out that *everyone* should accept a combinatorial solution, one that takes the new phenomenal characters of macroexperiences to be a matter of the combination

---

<sup>12</sup>One way to put the worry is that there are in principle reasons for thinking that functionalism cannot solve what Bourget (this volume) calls the "mapping problem." The worry mirrors underdetermination worries with functionalism about semantic properties; see, e.g., Kripke 1982, Bonjour 1998, Putnam 1977, Mendelovici and Bourget forthcoming, and Mendelovici 2018.

<sup>13</sup>See Gertler 2001 for a defense of phenomenal internalism.

<sup>14</sup>Another problem with the resulting view is that it makes the wrong predictions in certain cases, since the phenomenal characters of many phenomenal states do not match any items in the external environment (Bourget and Mendelovici 2014, Pautz 2006b, 2013b, and Mendelovici 2013, 2016, 2018, Chs. 3–4).

<sup>15</sup>Such a strategy is employed by several phenomenal intentionality theorists, who take some intentional states to be determined by phenomenal states while others are determined by their functional relations to phenomenal states. See Graham et al., 2007, Horgan and Graham 2009, Loar 2003, Bourget 2010, Pautz 2006a, 2013a, and Chalmers 2010, p. xxiv.



of other constituent phenomenal characters, had either by the macroexperience itself or by constituent experiences. Maybe the only way to get new phenomenal characters is out of combinations of old ones.

The upshot of this discussion is that the panpsychist's problem of explaining new simple phenomenal characters might be the same in kind as the problem of explaining changing quality spaces, a problem that everyone faces. While it might seem that the two problems admit of different solutions, I have suggested that the panpsychist can co-opt non-combinatorial solutions to the changing quality space problem and the changing quality space problem might have to be solved by appeal to mental combination anyways.

## 4 Implications for panpsychism

I have argued that panpsychism's combination problems are problems for everyone. This section considers the implications of this claim for objections to panpsychism based on the combination problem. I want to suggest that the fact that the combination problem is a problem for everyone suggests the *ignorance hypothesis*, on which we are ignorant of certain key facts about mental combination, similar to Stoljar's (2006) "ignorance hypothesis" used to defend (broad) physicalism. The ignorance hypothesis allows us to respond to two important objections to panpsychism based on the combination problem.

One objection to panpsychism based on the combination problem is that the combination problem undercuts one of the key motivations for panpsychism over physicalism, the argument from physicalism's perceived failure at offering an intelligible explanation of our experiences (see Strawson 2003). If the panpsychist cannot offer an intelligible explanation of our experiences either, then panpsychism is no better off than physicalism in this regard (see Goff 2009 and Carruthers and Schechter 2006).

The second objection is that the combination problem shows that panpsychism is false. If the facts about microexperiences and how they are combined do

not a priori entail the macroexperiential facts, then macroexperiences are not nothing over and above combinations of microexperiences, and panpsychism is false. Goff (2009) and Chalmers (2016) consider a conceivability argument against panpsychism along such lines, which is analogous to Chalmers' (1996) conceivability argument against physicalism.

If the panpsychist's combination problem is a problem for everyone, then this supports the ignorance hypothesis, which allows the panpsychist to respond to these objections. Everyone should agree that mental combination of the kinds the panpsychist requires *does* occur, so we know that there exists an intelligible explanation of mental combination, whether or not we do or can know it. This explanation might make reference to physical, functional, phenomenal, or other kinds of facts, or it might even take certain forms of mental combination to be primitive—for present purposes, it doesn't matter. But we don't currently have such an explanation. This suggests the ignorance hypothesis: we are ignorant of certain key facts about mental combination.<sup>16</sup>

The ignorance hypothesis allows us to respond to the second objection: We simply are not able to conclude that the facts about microexperiences and how they are combined do not a priori entail the macroexperiential facts. For all we know, the facts about mental combination that we are ignorant of secure the required entailment. So, conceivability arguments fail to show that panpsychism is false.<sup>17</sup>

---

<sup>16</sup>What of the alleged conceivability of panpsychist zombies, microexperiential and microphysical duplicates of human beings that lack macroexperiences (Goff 2009 and Chalmers 2016)? The panpsychist can accept the conceivability of panpsychist zombies, so long as she maintains that mental combination requires more than mere microexperiential and microphysical ingredients, such as “phenomenal bonding” relations or some such and claims that it is these further ingredients that we are ignorant of. The panpsychist might alternatively deny the conceivability of panpsychist zombies by maintaining that mental combination is a matter of some features of microexperiences that are not currently known by us.

<sup>17</sup>Such a response, in effect, casts doubt on the conceivability argument's premise conceivability premise, e.g., that it is conceivable for there to exist microexperiential zombies, understood as creatures having the same microexperiences combined in the same ways as the panpsychist stipulates are found in us but lacking macroexperience. (Goff's (2009) and Chalmers' (2016) arguments against panpsychism understand microexperiential zombies as having the same microexperiences (and sometimes physical properties) as us but not necessarily involving the same modes of combination. However, these alternative characterizations of panpsychist zombies would yield conceivability arguments only effective against versions of panpsychism on which the relevant modes of combination are entailed by the microexperiential (or perhaps physical) facts, and not versions that take mental combination to involve extra ingredients).

The first objection can also be avoided so long as the physicalist cannot similarly avail herself to an appeal to ignorance. If an appeal to ignorance is equally available to the physicalist and the panpsychist, then the panpsychist's intelligibility-based argument for panpsychism over physicalism still fails. I want to suggest that the panpsychist's ignorance hypothesis is more plausible than an analogous physicalist ignorance hypothesis: The classic arguments against physicalism (the conceivability argument, the knowledge argument, and explanatory gap worries) show not only that the physicalist has not offered an intelligible explanation of consciousness in terms of the physical, but, further, that there is no such explanation to be had. Given a certain conception of physical facts (e.g. Chalmers' (1996) conception as facts concerning the structure and dynamics of physical processes), we can see that no set of physical facts can a priori entail the phenomenal facts, and so, that not only do current physical theories fail to intelligibly explain consciousness, but so too would *any* other possible physicalist theories.<sup>18</sup> If this is right, then an appeal to ignorance cannot help the physicalist: We may be ignorant of many physical facts, but we know enough about what physical facts look like in order to see that they cannot result in phenomenal consciousness. In contrast, we have less of a clear idea of what a plausible account of mental combination might look like. As a result, we simply do not know that there is no possible account of mental combination that renders panpsychist explanations of macroexperiences intelligible. Our epistemic situation rules out a physicalist account of macroexperience but leaves open a panpsychist account.<sup>19</sup>

---

<sup>18</sup>This is, in effect, Chalmers' argument against physicalism in "Facing up to the problem of consciousness" (1995), and arguably the core reason to think that zombies, physical duplicates of us lacking consciousness, are conceivable. "But the structure and dynamics of physical processes yield only more structure and dynamics, so structures and functions are all we can expect these processes to explain. The facts about experience cannot be an automatic consequence of any physical account, as it is conceptually coherent that any given process could exist without experience." (p. 208)

<sup>19</sup>However, the very existence of the combination problem, and the related conceivability arguments by Goff (2009) and Chalmers (2016), suggest that mental combination cannot be explained in terms of a mere combination of physical, functional, or microexperiential ingredients. See also fn. 16.

## 5 Concluding remarks

I have argued that the panpsychist's combination problems are problems for everyone and suggested that this alleviates the panpsychist's worries concerning intelligibility. Before concluding, it is worth emphasizing that combination problems afflict our very understanding of the mind largely independently of any particular metaphysical theories of mind. These problems are pervasive and multi-faceted, arising for many different kinds of mental states and under many guises. And they are largely underappreciated. For example, much discussion of phenomenal unity focuses on simply characterizing the phenomenon rather than explaining it.<sup>20</sup> Similarly, much discussion of intentional structure focuses on determining rules for when simpler contents combine to form more complex contents rather than explaining how mental structure is possible at all.<sup>21</sup>

Given the pervasiveness and apparent intractability of combination problems, it is worth considering the possibility that we not only have not solved them, but that we simply cannot solve them. Perhaps we are “cognitively closed” (McGinn 1989) to them in that our minds simply cannot grasp how mental things can combine. It at least seems that we can intuitively understand items being spatially, causally, or temporally related in various ways, that we can understand them piling up, bumping each other around, and existing and changing through time (whether or not this is enough to understand physical combination). But mental combination arguably requires something more than that. It requires a new mode of interaction whereby mental things merge, blend, or otherwise become more than a spatiotemporally and causally integrated sum of their parts. Perhaps this is something we are simply not equipped to grasp, making the mind impossible for us to completely understand, and giving rise to an unbridgeable (by us) explanatory gap between mental combinations and their uncombined

---

<sup>20</sup>For instance, both Dainton (2000) and Bayne and Chalmers (2003) mainly aim to characterize phenomenal unity, rather than to offer an explanation of how it arises.

<sup>21</sup>For instance, both Dainton (2000) and Bayne and Chalmers (2003) mainly aim to characterize phenomenal unity, rather than to offer an explanation of how it arises.

parts that faces physicalists, dualists, and panpsychists alike.<sup>22</sup>

## References

- 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.
- Bayne, T. (2012). *The Unity of Consciousness*. Oxford University Press Uk.
- Bayne, T. J. and Chalmers, D. J. (2003). What is the unity of consciousness? In Cleeremans, A., editor, *The Unity of Consciousness*. Oxford University Press.
- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in the Philosophy of Science*, 9:261–325.
- BonJour, L. (1998). *In Defense of Pure Reason*. Cambridge University Press.
- Bourget, D. (2010). Consciousness is underived intentionality. *Noûs*, 44(1):32–58.
- Bourget, D. (forthcoming-). Anomalous panpsychism: A new approach to the mind-body problem. In Seager, W., editor, *The Handbook of Panpsychism*.
- Bourget, D. (MS). The underdetermination problem for conceptual role semantics and phenomenal functionalism.
- Bourget, D. and Mendelovici, A. (2014). Tracking representationalism. In Bailey, A., editor, *Philosophy of Mind: The Key Thinkers*, pages 209–235. Continuum.
- Bourget, D. and Mendelovici, A. (2016). Phenomenal intentionality. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Carruthers, P. and Schechter, E. (2006). Can panpsychism bridge the explanatory gap? *Journal of Consciousness Studies*, 13(10-11):32–39.
- Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press, Oxford.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3):200–19.
- Chalmers, D. J. (2010). *The Character of Consciousness*. Oxford University Press.
- Chalmers, D. J. (2016). The combination problem for panpsychism. In Jaskolla, L. and Bruntrup, G., editors, *Panpsychism*, pages 179–214. Oxford University Press.
- Coleman, S. (2012). Mental chemistry: Combination for panpsychists. *Dialectica*, 66(1):137–166.

---

<sup>22</sup>Many thanks to David Bourget and Luke Roelofs for helpful comments on earlier drafts of this paper.

- Coleman, S. (2013). The real combination problem: Panpsychism, micro-subjects, and emergence. *Erkenntnis*, (1):1–26.
- Dainton, B. F. (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. Routledge.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press, Cambridge.
- Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press, Cambridge, MA.
- Gaskin, R. (2008). *The Unity of the Proposition*. Number 2. Oxford University Press.
- Gertler, B. (2001). Introspecting phenomenal states. *Philosophy and Phenomenological Research*, 63(2):305–28.
- Goff, P. (2006). Experiences don’t sum. *Journal of Consciousness Studies*, 13(10-11):53–61.
- Goff, P. (2009). Why panpsychism doesn’t help us explain consciousness. *Dialectica*, 63(3):289–311.
- Goff, P. (2017). *Consciousness and Fundamental Reality*. Oxford University Press.
- Graham, G., Horgan, T. E., and Tienson, J. L. (2007). Consciousness and intentionality. In *The Blackwell Companion to Consciousness*, pages 468–484. Blackwell.
- Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4:31–52.
- Horgan, T. and Graham, G. (2009). Phenomenal intentionality and content determinacy. In Schantz, R., editor, *Prospects for Meaning*. De Gruyter, Amsterdam.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(April):127–136.
- James, W. (1890). *The Principles of Psychology*. Dover Publications, New York.
- Kriegel, U. (2011). *The Sources of Intentionality*. Oxford University Press.
- Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language*. Number 2. Harvard University Press.
- Leibniz, G. W. (1714/1989). The monadology. In Ariew, R. and Garber, D., editors, *G. W. Leibniz: Philosophical Essays*, pages 333–45. Hackett.
- Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly*, pages 356–61.
- Loar, B. (2003). Phenomenal intentionality as the basis of mental content. *Reflections and Replies: Essays on the Philosophy of Tyler Burge*.

- Lockwood, M. (1993). The grain problem. In Robinson, H. M., editor, *Objections to Physicalism*, pages 271–291. Oxford University Press.
- Maxwell, G. (1979). Rigid designators and mind-brain identity. *Minnesota Studies in the Philosophy of Science*, 9:9.
- McGinn, C. (1989). Can we solve the mind-body problem? *Mind*, 98(July):349–66.
- Mendelovici, A. (2013). Reliable misrepresentation and tracking theories of mental representation. *Philosophical Studies*, 165(2):421–443.
- Mendelovici, A. (2016). Why tracking theories should allow for clean cases of reliable misrepresentation. *Disputatio*, 8(42):57–92.
- Mendelovici, A. (2018). *The Phenomenal Basis of Intentionality*. Oxford University Press.
- Mendelovici, A. and Bourget, D. (forthcoming). Consciousness and intentionality. In Kriegel, U., editor, *Oxford Handbook of the Philosophy of Consciousness*. Oxford University Press.
- Morch, H. H. (2014). *Panpsychism and Causation: A New Argument and a Solution to the Combination Problem*. PhD thesis, Oslo.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4):435–450.
- Palmer, S. (1999). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences*, 22(6):923–943.
- Pautz, A. (2006a). Can the physicalist explain colour structure in terms of colour experience? *Australasian Journal of Philosophy*, 84(4):535–564.
- Pautz, A. (2006b). Sensory awareness is not a wide physical relation: an empirical argument against externalist intentionalism. *Noûs*, 40(2):205–240.
- Pautz, A. (2013a). Does phenomenology ground mental content? In Kriegel, U., editor, *Phenomenal Intentionality*, pages 194–234. Oxford.
- Pautz, A. (2013b). The real trouble for phenomenal externalists: New empirical evidence for a brain-based theory of consciousness. In Brown, R., editor, *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*, pages 237–298. Springer.
- Pitt, D. (2004). The phenomenology of cognition or what is it like to think that *P*? *Philosophy and Phenomenological Research*, 69(1):1–36.
- Putnam, H. (1977). Realism and reason. *Proceedings and Addresses of the American Philosophical Association*, 50(6):483–498.
- Revonsuo, A. (1999). Binding and the phenomenal unity of consciousness. *Consciousness and Cognition*, 8(2):173–85.
- Roelofs, L. (2014). Phenomenal blending and the palette problem. *Thought: A Journal of Philosophy*, 3(1):59–70.

- Roelofs, L. (2015). *Combining Minds: A Defence of the Possibility of Experiential Combination*. PhD thesis, University of Toronto.
- Roelofs, L. (2016). The unity of consciousness, within subjects and between subjects.
- Roelofs, L. (forthcoming). Can we sum subjects? evaluating panpsychism's hard problem. In Seager, W., editor, *The Handbook of Panpsychism*.
- Seager, W. (1995). Consciousness, information, and panpsychism. *Journal of Consciousness Studies*, 2(3):272–88.
- Seager, W. (2010). Panpsychism, aggregation and combinatorial infusion. *Mind and Matter*, 8(2):167–184.
- Seager, W. (2016). Panpsychist infusion. In Jaskolla, L. and Bruntrup, G., editors, *Panpsychism*, pages 229–248. Oxford University Press.
- Smythies, J. R. (1994). Requiem for the identity theory. *Inquiry*, 37(3):311–29.
- Stoljar, D. (2006). *Ignorance and Imagination: The Epistemic Origin of the Problem of Consciousness*. Oxford University Press USA.
- Strawson, G. (2003). Real materialism. In Antony, L. M. and Hornstein, N., editors, *Chomsky and His Critics*, pages 49–88. Blackwell.
- Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press, Cambridge.