

Ole Koksvik

## *Three Models of Phenomenal Unity*

**Abstract:** *There is something it is like for me to hear a seagull crying, something it is like to see a boat in the distance, and something it is like to suffer a slight headache. Each of these local conscious experiences have their own phenomenal character. The experiences are phenomenally unified just in case there is also something it is like to enjoy these and all the other local experiences I have at the relevant time together. For there is also something it is like to be me overall: my global conscious experience has a phenomenal character. But what is it like to be me overall? What is the relationship between the phenomenal characters of local experiences and the phenomenal character of the global experience to which they contribute? This paper argues that our concepts of local and global conscious experiences allow for three completely different conceptions of how the former combine into the latter. It also argues that this shows that our concepts of local conscious experiences, global conscious experiences, and of their relationship are much more permissive than we might have thought.*

### **1. Introduction**

At the beach, your feet being lapped by gentle, cool waves, you feel a mild breeze on your skin, and enjoy the warmth of the sun. You hear the sound of the ocean, a revving motorcycle, the buzz of voices engaged in conversation, and a seagull crying. You see the ocean, the reflection of sun on the sea, the gentle waves, people playing and swimming, and boats in the distance. Despite suffering a slight headache and being thirsty and slightly too hot, you feel relaxed and content.

Correspondence:  
Ole Koksvik, University of Bergen. Email: [ole.koksvik@gmail.com](mailto:ole.koksvik@gmail.com)

You may not stop to think about it, but there is *something it is like* to be you when all of this is taking place. You feel a certain particular way overall, at this and every other waking moment. When there is *something* it is like to be you, you are phenomenally conscious. *What* it is like varies greatly: think about your overall experience as you take a pleasantly hot shower after a long sleep, as you intently study a text, and as you hop around on one leg having stubbed your toe. The particular way it is like for you at a time is the *phenomenal character* of your overall experience at that time.

Just as your overall, or global, conscious experience has a certain phenomenal character, so too do all your *local* conscious experiences. For example, there is something it is like to have a dull headache — a dull headache has a certain phenomenal character — and so does feeling a mild breeze on one's skin, hearing a buzz of background conversation, and so on.

So: the world we live in contains certain entities — local and global conscious experiences — which have certain properties — phenomenal characters. These facts are not questioned here. Instead the aim is to consider aspects of the relationship between these properties and entities, given that they exist.

One aspect of the relationship between local and global experiences which has received a great deal of recent interest is usually called 'the *unity of consciousness*' (see, for example, Cleeremans and Frith, 2003; Bennett and Hill, forthcoming). There are several ways in which conscious experience might be unified (Bayne and Chalmers, 2003; Bayne, 2010), but particular interest attaches to *phenomenal* unity. There is something it is like to hear a seagull crying, and something else it is like to feel a mild headache. These local experiences are phenomenally unified just in case there is also something it is like to enjoy these and all others *together*.

In humans the norm seems to be that all local experiences a person enjoys at one time are phenomenally unified. It is natural to wonder whether it *has* to be this way. Is there a natural law which guarantees this? Or, are the ideas that one and the same person could have more than one overall conscious experience at a time, or that a local conscious experience might fail to belong to a global experience, perhaps even incoherent, so that the unity of conscious experience holds with stronger than nomic necessity (Bayne and Chalmers, 2003)? These are important questions, since their answers are likely to shed light on the very natures of conscious experiences, subjects, and more besides (as Bayne and Chalmers note). There are also other interesting and important questions. What are the necessary and sufficient conditions

for conscious experience to be unified (one of the questions Watzl considers in this symposium). In virtue of what is consciousness unified when it is (Masrour, forthcoming)? And so on.

In this paper I set these questions to one side in order to focus on an aspect of the relationship between local and global experiences which so far has largely been overlooked.<sup>1</sup> A set of local experiences are phenomenally unified just in case there is something it is like to enjoy them all together. But *what* is it like?

Some answers to this question are ruled out from the outset, because we know that the phenomenal character of a global experience somehow reflects the phenomenal characters of the local experiences which contribute to it. Answers which fail to deliver this result are simply not contenders. We want to know how the phenomenal character of the global experience reflects the phenomenal characters of the local experiences, or, as we might put it, how the phenomenal characters of the local experiences *combine* into the character of the global experience. Thus we get the Combination Question about Conscious Experiences:

**The Combination Question:** How do local conscious experiences combine into a global conscious experience?

The aim of this paper is to argue that several *radically different* models of how combination takes place might be true. Put differently: consideration of the Combination Question reveals that our concepts of local and global experiences are surprisingly permissive, in that they allow that local and global experiences might turn out to be a range of very different kinds of things.

The route I take to these conclusions is as follows. In §3 I present three models for how local experiences might combine into a global one. Then I argue that each of the models merits being taken seriously (§4), and, moreover, that each is compatible with all central platitudes about conscious experiences (§§5–6). Finally I draw out the consequences for our concepts of local and global experiences, and of the

[1] Notable exceptions include Shaw (no date), Pautz (2013), and Dainton (2008; 2010), see also Lee (forthcoming). The first two authors delineate the limits for which local experiences can combine into global ones, and rule some combinations out as impossible. They thus come close to the subject matter discussed in this paper, but do not actually enter into the discussion I want to raise, and seem in fact to presuppose the truth of the no-context-dependence model (§3 below). Dainton's discussion comes closer, especially in his 2008, §§9.2–9.4, but the issues he focuses on and those discussed here cross-cut. To illustrate: a central question for Dainton is whether any experience has 'a character that it *could not have had* in the absence of [a certain other] experience' (*ibid.*, p. 271). This might be true on at least two of the three models I outline, but might also be false on *all* of them. See also n. 19 below.

relationship between them (§7). Before getting this far, however, we need to play some defence.

## 2. Coexistence

Some might worry that there is tension between the Combination Question and nearby views about conscious experiences. In this section I argue that the Combination Question can coexist peacefully with a number of nearby positions, in the sense that the positions' turning out to be true (or false) would not show the question to be *misguided*, *uninteresting*, or *unimportant*.

A position's being true might show that a question is misguided, in the sense that it embodies a false presupposition. To illustrate, many hold that the properties which determine what it is like to enjoy a conscious experience are the very same as those which determine what the experience represents. On this assumption it is misguided to ask which of the two types of properties is more fundamental. There is no distinction between the 'two', so neither is more fundamental than the other.

A position might also show a question to be uninteresting. For example, it might show that the answer is trivial for all the cases we care about. Suppose it were convincingly shown that in order to have rights a being must have self-consciousness, and also that no non-human animal on earth has this capacity. Against this background the question 'what rights do animals of type x have' would be uninteresting since for all values of x about which we care the answer would be: 'none'.

Finally, a position might show a question to be unimportant. A thoroughgoing epiphenomenalism might do so in the present case. If conscious experiences are 'nomological danglers' (Smart, 1959), with no impact on other aspects of the mental life of a person — her beliefs, desires, emotions, rationality, or anything else — one might be forgiven for thinking that the Combination Question is unimportant. It would also be unimportant if it turned out to be theoretically isolated — if answering it a certain way had no tendency to put pressure on other parts of theory — so any position which indicated this would also threaten the Combination Question.

### 2.1. Unity

Let us begin with phenomenal unity. Bayne and Chalmers (2003) define it this way:

**Phenomenal Unity Thesis (PUT):** Necessarily, for any conscious subject of experience ( $S$ ) and any time ( $t$ ), the simultaneous conscious states that  $S$  has at  $t$  will be subsumed by a single conscious state — the subject's total conscious state.

PUT is obviously closely related to the Combination Question, so if any position could not easily coexist with the question, one might have thought this would be it. However, the Combination Question can coexist just fine with PUT.

If PUT is true, then (in our parlance) all local experiences had by a subject at a time necessarily combine into a global experience. This clearly does not show the Combination Question to be misguided: if local experiences always combine to a global one it makes good sense to ask how they combine. Accepting PUT also has no tendency to show the Combination Question to be uninteresting; that local experiences always combine to a global one does not guarantee that they always combine the same way. Nor does PUT indicate that the Combination Question is likely to be theoretically isolated; if anything, it indicates the opposite.

The Combination Question is also not threatened by PUT's being false. The debate about the unity of experience asks whether local experiences always combine into a global one, and if so, with what kind of necessity. No one doubts that *some* local experiences *sometimes* combine into a global experience, so if PUT is false we can still ask how local experiences combine into global ones, when they do. The other reasoning goes through in parallel fashion; the Combination Question embodies the presupposition that local experiences sometimes combine to global ones, but the falsity of PUT entails no denial of this claim, nor does it entail or indicate that the combination will always take place the same way, nor that the question is theoretically isolated.

## 2.2. *The One Experience View*

Another position that might be thought to put pressure on the Combination Question is the 'one experience view' of Michael Tye (2003). He argues that what should be regarded as an experience is a person's entire stream of consciousness from the time she wakes up from dreamless sleep until the time she falls into such sleep again (or until consciousness is otherwise interrupted).

I think it is clear that, at least in this context, the issues raised by this view are merely verbal (Chalmers, 2011). It does not matter whether

we reserve the label ‘experience’ for the streams of consciousness from sleep to sleep, or whether we more liberally apply it also to tastings of honey, smellings of roses, and stubbings of toes, on the one hand, and the overall experiences had at various times, on the other. For even in the former case experience obviously has distinguishable *parts* to it, some of which are unified in a way others are not (Dainton, 2008, p. 72). For example, my drowsiness and my perceptual experience as of faint light through the curtains as I first wake up combine in a way that my drowsiness does not combine with the taste of my dinner. We can label this what we will, but the facts on the ground remain the same.

### 2.3. *Atomism and Holism*

The Combination Question also coexists peacefully with both *atomism* and *holism* about conscious experiences. Tim Bayne writes:

Theorists who adopt an atomistic orientation assume that the phenomenal field is composed of ‘atoms of consciousness’—states that are independently conscious. Holists, by contrast, hold that the components of the phenomenal field are conscious only as the components of that field. Holists deny that there are any independent conscious states that need to be bound together to form a phenomenal field. Holists can allow that the phenomenal field can be formally decomposed into discrete experiences, but they will deny that these elements are independent atoms or units of consciousness. (Bayne, 2010, p. 225)

I see two salient ways of understanding holism: as a *fundamentality* claim, or as a *modal* claim (see Schaffer, 2010; and also Lee, forthcoming; Dainton, 2008; 2010). Understood the first way, holism says that the global experience is more fundamental than and grounds the local experiences. Understood the second way holism says that the local conscious experiences would not be conscious if they were not components of a global experience.

Correspondingly, there are two ways to understand atomism. As a claim about fundamentality it is the claim that local conscious experiences are more fundamental than, and ground, global experience. As a claim about modality it says that, even though a given local experience typically occurs as a component of a global experience, maybe even in all cases with which we are acquainted, and perhaps even by nomological necessity, the local experience could have existed outside of it.

The Combination Question coexists peacefully with both atomism and holism because it does not itself concern fundamentality or modality. The Combination Question is not about whether global or

local experiences are more fundamental. It is instead about how the phenomenal characters of local experiences relate to that of the global experience to which they contribute. Similarly, it is not a question about the modal profile of local experiences, and in particular not about whether local experiences can occur alone. Instead it is a question about the cases where they do not.

Neither atomism nor holism reveals the Combination Question to be misguided. If atomism read as a modal claim is true, local experiences can occur in isolation, but this does not show that they cannot contribute to the character of a global experience. Evidently they often do, and the Combination Question asks about those cases. Atomism understood as the claim that local experiences grounded, and are more fundamental than, global experiences, also does not show that the Combination Question is incoherent: local experiences can contribute to a global experience regardless of which is more fundamental (and also if both are equally fundamental). This reasoning also shows that holism on the fundamentality interpretation does not reveal the Combination Question to be misguided.<sup>2</sup>

Perhaps the most challenging view for the Combination Question is holism on the modality interpretation, but even this view coexists peacefully with the Combination Question. On the modality interpretation holism says that local conscious experiences could not exist on their own: there is something it is like to have a headache only if it occurs as part of a global experience. But we may simply think of this as a condition for what it takes for a local conscious experience to contribute to a global one. Once that condition is met, the question still remains of *how* the local experience contributes, when it does. That is the Combination Question. So, atomism and holism, on either interpretation, fail to reveal that the Combination Question is misguided.<sup>3</sup>

In summary, the Combination Question coexists peacefully with both unity and its negation, with the one experience view, and with holism and atomism.

---

[2] I am talking here about fundamentality of objects, in the manner of Schaffer (2010).

[3] Holism on the modality interpretation might be thought to show that, in some sense, there are no local experiences. In §§6 and 7 below I argue that it instead shows that local experiences are very different sorts of entities that what we originally thought.

### 3. Three Models of Combination

In this section I describe three very different models of combination.<sup>4</sup> At this point getting the general ideas across is more important than the details, so the description is given in broad brushstrokes, with significant reliance on metaphor. In subsequent sections we shall make many aspects of the models more precise.

Colour plays an important part in the description. It is useful, since we are familiar with colours resembling one another along several distinct dimensions. In the models, colour stands for phenomenal *feel*, that is, phenomenal character. We can think of the phenomenal characters of overall experiences as a value in a highly complex multi-dimensional space, about which dimensions of variation we as yet know relatively little, and — this being one way to characterize the point of the present article — such that the interplay between those dimensions is even more unknown.

#### 3.1. The No-Context-Dependence View

What I take to be the dominant view is nearly always implicitly rather than reflectively endorsed. This view holds that the character of overall or global experience results from simple ‘addition’ of the characters of individual experiences. (Frank Jackson is one who endorses this view in conversation, though not thus far in print.)

On this *no-context-dependence* view we can liken overall conscious experience to a *mosaic*, onto which small coloured tiles are fitted. Different local conscious states are represented by tiles of different colours, and which colour is placed on the mosaic depends only on which contributor states are instantiated. For example, having a certain visual perceptual experience causes a red tile to go on the mosaic; an emotional experience, a yellow tile, etc.

On this view, the character of my overall conscious experience is guaranteed to be *exactly similar in a respect* each time I am in the same local conscious state. For example, each time I have the same emotional experience, the mosaic will contain a region of the relevant colour. If we wish, we can imagine that the size of the tile depends on the intensity of the experience, and on whether or not it is attended to. Further embellishments could be added. However, the main point is

---

[4] Strictly speaking we shall be concerned with model *types*, within which there is significant scope for variation. For ease of presentation I will bracket this in the main text.



that a particular local conscious experience always contributes a tile of the same colour to the mosaic.<sup>5</sup>

### 3.2. *The Weak-Context-Dependence View*

We might instead imagine each contributor state being associated with a colour *determinable*. Which determinate shade of that colour is placed on the mosaic would, according to this idea, depend on which other tiles are already present there. This corresponds to the thought that the contribution a local experience makes to the phenomenal character of the person's overall conscious experience might be similar each time the contributor state obtains, while still depending on which other experiences the person is undergoing at the time.

On this *weak-context-dependence* view there is no guarantee of *exact* similarity in any aspect of my overall experience each time I am in the same local conscious state. For example, each time I suffer a dull headache, there is no guarantee that what it is like to be me overall will be exactly similar in any respect. But there will be *approximate* similarity in a respect between such instances; the same colour determinable is present each time. The character of a person's overall experience at a time is on this view a more complex function (than straightforward addition) from the characters of local contributor experiences.

### 3.3. *The Strong-Context-Dependence View*

Finally, our overall experience may not be like a mosaic at all, but instead more like a shallow pool of water, into which a vial of a certain liquid is poured whenever the subject has a certain local experience. The liquid colours the water through a chemical reaction. Different liquids have different effects, and the effect of adding a liquid depends on the chemical composition of what is in the pool. Adding the same amount of the same liquid can yield different results on different occasions.

On this *strong-context-dependence* view we are not even guaranteed that there will be similarity in which colour determinable is instantiated each time I have a particular local experience. Even though the same amount of the same type of liquid has been poured in, the overall result with respect to colour may be completely different (we are stipulating): there need be no region of the pool which even

---

[5] Here and throughout we bracket the complication that a colour looks different against differently coloured backgrounds; metaphors do not require full fidelity to reality to be evocative and useful.

has the same colour determinable in the two instances. On this model, the impact of having a particular conscious experience can depend on which other experiences the person is currently undergoing to such a degree that there is no guarantee of *even approximate similarity in a respect* in the resultant overall phenomenology. Even on this view, however, the phenomenal character of the overall experience results from a function from the characters of local contributor experiences, albeit an extremely complex one.

#### 4. Taking the Models Seriously

We are intimately familiar with our own conscious experiences, so one might have expected us to easily decide between these very different models. Quite surprisingly, we cannot easily do so, or so, at least, I shall argue.

We begin by noting some reasons to take the models seriously.

The no-context-dependence model is pretty clearly the one that most closely aligns with our pre-theoretic ideas about local and global conscious experiences. (This helps to explain why it is usually simply presupposed.) It is attractive because it gives the simplest picture of the relationship between local and global conscious experiences, making that relationship easy to understand. This is arguably a virtue of any model, and in particular of a model that purports to be about something we know as well as we do our conscious experiences.

Secondly, philosophical conservatism may be thought to speak in favour of the no-context-dependence view. A number of positions have a stake in the phenomenal characters of local and global conscious experiences — representationalism and dogmatism, for example — and yet discussions of the relationship between local and global conscious experiences are absent from the defences of these views. That seems to indicate that the relationship must be thought to be simple and straightforward, which is what this model delivers. Inasmuch as the views can be shown to *require* the truth of the model,<sup>6</sup> and inasmuch as the views are otherwise attractive, we get further motivation to take the no-context-dependence model seriously.

Finally, direct reflection on experience gives a further reason to take this model seriously. It certainly seems *prima facie* plausible that *the very auditory experience* I am currently enjoying might just as easily have occurred even if my visual perceptual experience had been completely different than it actually is (Dainton, 2010, pp. 273–5;

[6] I think strong arguments can be mounted in both cases, but I must leave this for another occasion.

Shaw, no date, p. 10; Pautz, 2013, §3.2.).<sup>7</sup> And it is natural to think that the model which captures this is the no-context-dependence model.

On the other hand, the two context-dependence models have virtues of their own. One reason to take context-dependence seriously results from reflection on the interaction between moods, especially sadness and elation, and other conscious experiences. As most of us have experienced first-hand, the world seems a very different place when one is elated than when one is profoundly miserable. Even going for a familiar walk can be a very different experience in a buoyant mood. It is as if the good feeling bleeds into every other experience: even mundane things *look* sparkly and full of promise at those times, utterly different from their grey and hopeless appearance in negative moods.

There is also experimental evidence pointing in the same direction. It is possible to set things up so that one can control the volume and frequency-distributions of sounds that people normally hear unaltered in connection with everyday activities such as eating or rubbing one's hands together.<sup>8</sup> Using this experimental paradigm, it has been shown that boosting the volume of the sound a subject hears, or boosting the proportion of high-frequency sounds, affects the degree to which potato chips are perceived as crisp or stale when eaten (Zampini and Spence, 2004), and the perceived roughness or smoothness of one's palms when the hands are rubbed together (see Jousmäki and Hari, 1998). The data of course lends itself to multiple interpretations, but a salient one is that the phenomenal characters which the eating of crisps and rubbing of hands contribute to the character of the subject's overall experience depend on the phenomenal context in which they occur, that is, on which other phenomenal experiences the person is undergoing.<sup>9</sup>

Both of these considerations point in the direction of context-dependence of some sort, but it is perhaps not clear that they support anything stronger than the weak version of that thesis. However, I

- 
- [7] James R. Shaw puts it this way: '[T]he conscious combinatorics of most modalities seems to be boring... [in that] [i]t seems that all sorts of... permutations... are all clearly conceivable, and hence (barring special explanations) possible. Color experiences seem like they can be arbitrarily permuted within a single visual field, and across visual fields... The same goes for the sensory modalities of gustatory, auditory, and tactile phenomenology...' (Shaw, no date, p. 10).
- [8] The subject wears a headset, and receives the sounds she would normally receive directly through the air, through it instead, manipulated in various ways, or not.
- [9] In a different experimental paradigm which points in the same direction, visual motion has been found to affect tactile motion, and *vice versa*. See Konkle *et al.* (2009). For philosophical discussion of all these cases, see Dainton (2010).

think we can find evidence supporting even strong-context-dependence in our own experience.

Consider the contribution experience from a particular sensory modality can make to the character of one's overall experience.<sup>10</sup> In normal cases, visual perceptual experience resolves into distinct objects, represented as having various properties and as standing in various relations to each other. At least in my own case, however, there are phenomenal contexts in which visual perceptual experience seems to contribute nothing but a certain 'visualness' to the character of my overall experience. Similarly with auditory perceptual experience; in some settings all that is contributed is 'auditoriness'; the experience of being exposed to some sound or other. Again, these phenomena lend themselves to multiple interpretations, but a salient one is that perceptual experience can make extremely different contributions to the character of a person's overall conscious experience in different phenomenal contexts; so different as to not guarantee resultant similarity in any respect.<sup>11</sup>

It seems, then, that we have reason to take each of the models seriously at the outset, in the sense of regarding each as a serious contender for truth and assigning it non-negligible credence. In the next section I discuss the central platitudes about local and global conscious experiences, before arguing, in §6, that each of the models can account for all of them. I take this to show that the epistemic situation we find ourselves in with respect to the three models is at least somewhat *robust*: we should still take all three models seriously even after conceptual investigation. This does not show, of course, that we cannot ultimately decide between the models. But it does show that the considerations which do end up deciding between them will either be empirical, or, if conceptual, subtle and unobvious.

## 5. Platitudes

What are the central platitudes about local and global experiences that are common knowledge between us?<sup>12</sup> We have already begun answering this question, by noting that *there are* both local and global

[10] Perhaps there strictly speaking are no sensory modalities, or perhaps there are many more than what we are often taught (Macpherson, 2011). Nothing here hinges on this.

[11] In these examples, 'visualness' and 'auditoriness' are contributed throughout, but these are too general properties to count as similarity in a respect in the relevant sense. (Specifically, they are attitude-specific, and not content-specific, aspects of phenomenal character, see Koksvik, 2011, §4.1.)

[12] They are propositions we all know about local and global conscious experiences, that we all know that we all know, and so on; see Lewis (1972).

experiences, both of which have phenomenal characters.<sup>13</sup> This is the most basic of the platitudes.

When you hop around on one leg having just stubbed your toe there is some particular way it is like for you overall. That is the phenomenal character of your overall conscious experience. You are unlikely to be able to fully describe that character: our language here is poorly developed. Moreover, you yourself may not even know all there is to know about the experience. None of this throws doubt on the claim. The claim is an ontological one, about what there is in the world. With respect to this claim, our descriptive or epistemic abilities are neither here nor there.

All local conscious experiences have phenomenal characters too: there is something it is like to taste a fresh apple, for example. There is no commitment here to local experiences being *simple*: tasting an apple may still be a highly composite affair, perhaps dependent on the operation of the olfactory system, for example (Smith, 2012). Similarly, the painful experience of stubbing one's toe may have parts or aspects to it (throbbing, pressure, heat, ...). Again, none of this endangers the platitude: if what it is like to taste an apple (or: to enjoy the flavour of the apple) is a composite affair, there is still something it is like to do so.

There are local and global conscious experience with phenomenal characters. And we can add: the phenomenal character of a global conscious experience somehow reflects the phenomenal character of the local conscious experiences which contribute to it. That is a second platitude. We want to know *how*. That is what the models are for.

So far we have been concerned with ontology. It is time to add a platitude about epistemology. The third platitude about conscious experiences says that we all know a fair bit about what it is like to undergo a broad range of conscious experiences. 'A fair bit' and 'a broad range' are vague expressions, so this platitude itself is vague. Its truth is compatible with big gaps in our knowledge. That is a good thing, since in some cases there may be a lot we don't know.

Connoisseurs of fine food and drink often produce descriptions of the associated experiences that to the uninitiated seem implausibly complex. (No doubt there are cases where the aim is merely to appear sophisticated, but let us set those to one side.) It may be that experiences change as one becomes more familiar with them. But an experienced taster might also notice aspects of the experience that were

---

[13] To avoid the merely verbal dispute discussed in §2.1 above we may have to reword this claim.

there all along, but which the novice overlooked. If the latter is true, and if the point generalizes, we often don't know exactly what the characters of even local experiences are: aspects — even significant aspects — elude us. All of this is perfectly compatible with the second platitude. The claim is the modest one that we all know *a fair bit* about what it is like to undergo *a broad range* of experiences, not the overly ambitious one that we know exactly what it is like to undergo all experiences.<sup>14</sup>

Epistemic modesty is a virtue, but we should not throw the platitude-baby out with the modesty-bathwater. If the second platitude were false, many of our everyday practices would be incomprehensible.

For example, if I did not know a fair bit about what it is like to suffer a headache I would lack reason to sympathize with you when you suffer one.<sup>15</sup> Similarly, I would seem to lack reason to envy you the taste of a rare and fine fruit (or wine or cognac), or to attempt to avoid humiliating situations, unless I knew a bit about what it is like to undergo the relevant experiences. So, that we all know a fair bit about the character of a range of experiences is overwhelmingly likely.

We need to add one final claim to our collection, namely that we successfully refer to the phenomenal characters of our conscious experiences often enough for discussions such as this one to proceed sensibly.<sup>16</sup> This point is obviously related to the others; that there is a character (that there are properties) to refer to, and that we often enough know well enough what that character is, are both plausible general preconditions for successful reference.

There are additional problems with phenomenal character, stemming from the fact that we do not have direct access to other people's experiences, but it is plausible that these challenges can be and routinely are overcome. Again, our everyday practices would be hard to make sense of otherwise: we not only feel but also *express* sympathy, regret, aversion, and so on, and we have no reason to think that our communicative practices here are defect.

The central platitudes about conscious experiences are thus claims about the way things are with respect to metaphysics, epistemology, and reference. As starting points for theorizing these are about as safe as they come, and I will defend them no further here.

[14] For a sceptical view, see Schwitzgebel (2008).

[15] At least, I would lack a reason which it seems obvious that in fact I have. I might not lack reason entirely; perhaps I could rely on knowledge of other experiences and a solid dose of extrapolation. Thanks to Mette Kristine Hansen for discussion here.

[16] Thanks to Barry Smith for spirited discussion here.

## 6. Accounting for the Platitudes

The platitudes capture, I take it, the central features of conscious experience which are common knowledge among us. One might have thought that this would suffice to decide between our three models, given how different they are from each other. However, I will now argue that all three models can account for each of the platitudes: each model is capable of giving content to the platitudes in ways that retain their plausibility.<sup>17</sup>

### 6.1. First Platitude

Let us begin with the claim that local and global experiences have phenomenal characters. It is easy to see that each model allows for a straightforward account of the characters of global experiences. In our metaphors, colour stands for phenomenal feel. On the first two models, the colour-pattern of the mosaic is the analogue for the phenomenal character of the overall conscious experience, and on the third, the colour pattern of the pond is.

Can the models also account for local conscious experiences having phenomenal characters? On the no-context-dependence model, each time a person has a certain conscious experience, a tile of a certain colour is placed on the mosaic. The colour of the tile is a natural analogue of the phenomenal character of that local experience.

On weak-context-dependence, each time a person undergoes a certain conscious experience, a tile of a certain colour determinable is placed on the mosaic, and which determinate colour is placed there depends on which other conscious experiences the person is enjoying — in the metaphor, on which other tiles are already on the mosaic. Suppose that the contribution a headache makes is always a shade of yellow. The natural suggestion is that the phenomenal character of the local conscious experience is what those tiles have in common: their ‘yellowishness’. Again, a simple idea, and it gives plausible content to — an unproblematic way of accounting for — the claim that local experiences have characters.

On strong-context-dependence our overall conscious experience is like a shallow pool, into which a vial of a certain liquid is poured each time a person enjoys a certain local experience. It might seem difficult to see how local experiences can have phenomenal characters on this model. After all, the liquid poured into the pond need not itself be coloured.

[17] I set the platitude about reference to one side here, and concentrate on showing that each of the models account well for the first three.

As it turns out, however, there are at least two relatively straightforward ways of giving content to the first platitude, even on strong-context-dependence. The view is compatible with no similarity in phenomenal character of overall experience between two occasions when the subject enjoys a certain local experience. However, the model is also compatible with there being a range of cases where such similarity *does* obtain. For example, we can imagine a range of ‘uncluttered’ phenomenal contexts, where the character of overall experience results from a small number of not very intense contributors. Suppose that when a person acquires a headache in such a context the pond invariably takes on a yellowish hue in some region. Then, the phenomenal character of this local experience could be equated with yellowishness, just as on weak-context-dependence.

There is another salient possibility: we might identify the character of a local conscious experience with the *pattern of change* in the character of overall experience each time the local experience contributes to it. To get the idea, consider the volume of some music played over a stereo being gradually doubled over a five-second period. Now imagine this change being carried out on many different occasions, with very different music being played. Obviously, there need be no similarity in what it is like to hear the music after the change between each of these instances. Equally obviously, there is something it is like to experience the volume being doubled in this way. Now imagine a visual scene, which first is under very low illumination, but then, over a period of five seconds, comes to be illuminated by intense flood lighting. It is clear that there need be no similarity in what it is like to perceive the scene after the change, since what that it is like will depend on the scene. Still, there is something it is like to perceive illumination being increased in this way.

Analogously, even if there is no similarity, on different occasions, in what it is like to have an overall experience to which a certain local experience contributes, there may still be something it is like to have that local experience if there is a pattern of change in what it is like before the experience begins to contribute, to after. We get what we might call *the difference view* of the character of local conscious experiences, which identifies the phenomenal character of local conscious experiences with the pattern of difference it makes in the phenomenal characters of global conscious experiences to which it contributes:

Difference view:  $PC(L) = \Delta PC(G)$

Even on the strong-context-dependence model, then, we can give content to the first platitude.



### 6.2. *Second Platitude*

The second platitude is that the phenomenal character of the overall conscious experience somehow reflects the phenomenal characters of the local experiences. Unsurprisingly, each of the models account for this fact, since that is precisely what they are designed to do.

### 6.3. *Third Platitude*

The third platitude says that we all know a fair bit about what it is like to undergo a broad range of conscious experiences. Can the models all account for this platitude?

It is important to be clear on what the question is here. I am the first to admit the difficulty of the question: ‘how do we know what the characters of our experiences actually are?’<sup>18</sup> However, here our task is not to answer *that* question, but rather to consider whether any *new* problems are introduced by the models. If not, the models give content to the platitudes in ways that *preserve* their plausibility, since it gets *no harder* to see how the platitudes might be true, given the models.

We have mentioned above the possibilities that global conscious experiences are more fundamental than local ones, and *vice versa*, metaphysically speaking. Similarly one might wonder whether either local or global conscious experience are epistemically prior to the other. It might be that the route to knowledge about global conscious experiences goes through knowledge of local conscious experiences. It might be that the route goes the other way. It might be that it goes sometimes one way and sometimes the other, and it might be that knowledge of one is obtained independently of knowledge of the other. In what follows I try to show that, either way, all of the models preserve the plausibility of the third platitude.

#### *No-context-dependence — local experiences are epistemically prior*

If local experiences are epistemically prior, I take it that a model could only fail to preserve the plausibility of the third platitude with respect to knowledge of the phenomenal characters of local conscious experiences if the very picture it painted of local experiences somehow made it incomprehensible that we could know about their properties. But on no-context-dependence, the phenomenal character of local conscious experiences correspond to the colour of the relevant tiles, so it seems clear that no such problem arises here.

[18] I argue elsewhere (Koksvik, forthcoming) that this is rather more difficult than some recent literature will have it.

It is also relatively easy to see how we could know about the characters of our global experiences, since on this model the characters of global experiences result from simple addition of the characters of local ones. Coming to know about the character of global experiences would then be a matter of using one's knowledge of local experience to come to know that one's global experience at least partly has such-and-such a character, namely the character of the local experience. This does not settle whether we know all there is to know about the characters of our global experiences; we may 'overlook' many local experiences in this process of coming to know the character of global experience. But we will at least be in a position to know a fair amount about its character, and certainly enough to preserve the plausibility of the platitude.

*No-context-dependence — global experiences are epistemically prior*

If global experiences are epistemically prior, I take it, as before, that a model can only fail to preserve the plausibility of the third platitude with respect to knowledge of the phenomenal characters of these experiences if the picture it paints of them somehow makes it incomprehensible that we could know about their properties. But on the no-context-dependence view, the phenomenal character of global conscious experiences is represented by the pattern of colours on the mosaic, so it seems clear that no problem of this nature arises here.

We can also make good sense of the process of coming to know the characters of local conscious experiences here. The same local experience contributing guarantees that there will be exact similarity in some respect of the overall experiences to which it contributes, so that process can be understood as recognizing the aspect of exact similarity between different instances. Such a capacity is not at all mysterious.

*Weak-context-dependence — global experiences are epistemically prior*

Let us begin consideration of how weak-context-dependence fares on the assumption that global experiences are epistemically prior. To begin with, just as on no-context-dependence, the phenomenal character of global conscious experiences correspond to the patterns of colour on the mosaic, so no special problem is introduced with respect to knowledge of that character here.

Knowledge of the phenomenal characters of local conscious experiences can be accounted for in much the same fashion as before. It is

true that the same local experience contributing does not guarantee that there will be *exact* similarity in any respect of overall experience. But it does guarantee *approximate* similarity (in a respect). If knowledge of the phenomenal characters of global experiences is epistemically prior we could come to know about the characters of local experiences through a process of *abstraction*: from a number of cases of a headache being present we would abstract away from minor differences and learn what they each had in common — their yellowishness. So, if global experiences are epistemically prior, the third platitude is well accounted for on the weak-context-dependence model.

*Weak-context-dependence — local experiences are epistemically prior*

If local conscious experience are epistemically prior, I assume as before that the model only fails to preserve the plausibility of the platitude with respect to the possibility of knowing about the characters of local experiences if something about the very picture it paints of what local experiences are makes it incomprehensible that we could know about their properties. On weak-context-dependence, local experiences are united by phenomenal similarity of a certain sort, what we have described through analogy of the yellowishness that all shades of yellow have in common. It is very plausible that we can know about such similarity in our experiences, so no new problems arise here.

For our knowledge of the phenomenal characters of global experiences there are two possibilities. First, it may be that we only know about the characters of our overall experiences at a relatively modest level of precision: we know that the overall experience has some shade of yellowishness on it, but we do not know which shade. Another possibility is that we learn rules for how other experiences ‘pick out’ a particular shade of the relevant colour, allowing us to gain exact — or at any rate more exact — knowledge about what an aspect of the overall experience will be like, each time a particular contributor state is present. Again, this platitude is modest, so both options are adequate.

*Strong-context-dependence — global experiences are epistemically prior*

On strong-context-dependence, the phenomenal character of global experiences corresponds to the colour on the surface of the pond. Applying the same assumption as before it seems clear that no new epistemic problems arise for knowledge about those phenomenal

characters: there is no special difficulty in understanding how we could know about such properties.

Strong-context-dependence is, we noted, compatible with no similarity in any respect of phenomenal character between two instances of a local experience occurring, but also compatible with such similarity arising in a range of ‘uncluttered’ contexts. On the development of the strong-context-dependence model which equates the phenomenal character of a local experience with the colour that emerges in such contexts, the epistemic situation with respect to the characters of local experiences is parallel to that which we saw under weak-context-dependence: we can come to know about the characters of local conscious experiences through a process of abstraction. We need only add that we recognize the relevant uncluttered contexts, and use them to acquire the relevant knowledge.

On the other development of the model considered above we arrived at the difference view, on which  $PC(L) = \Delta PC(G)$ . Here, the story is much the same as on weak-context-dependence: we would come to know about the characters of local experiences through a process of abstraction. This time, however, we abstract away from the differences between individual instances, not to arrive at a similarity like yellowishness, but instead to arrive at the pattern of similarity in *changes* in characters of global experiences. Of course, such patterns would not be simple, or easy to describe verbally. But all of this is unproblematic, since this is true of phenomenal characters in general.

We might initially lock on to a pattern of change in experiences we have actually enjoyed, but still end up knowing something that exceeds this. Knowing the character of a local conscious experience might amount to possessing an amount of counterfactual information regarding how ‘adding’ this local experience would change a range of possible global experiences: to know what it is like to smell a rose is, on this view, to know how my global experience would change if I did.<sup>19</sup>

*Strong-context-dependence — local experiences are epistemically prior*

If local experiences are epistemically prior, the picture looks a little different. On the development where local experiences result in the pond being tinged with a certain colour in uncluttered contexts we can still know *something* about the character of our overall experiences if local experiences are epistemically prior: it will be partly yellowish in

[19] Special thanks to Rachael Briggs for helpful discussion here, and to Sharon Berry for the idea.

those contexts. But here it seems that what is salvaged is not enough; we seem able to know about the characters of our global experiences even when things get busy.

On the other development, the idea of local experiences being epistemically prior does not even seem to get off the ground. There is no in principle problem with knowing about entities of this kind, that is, with knowing about capacities to change other entities. But we cannot make sense of knowing about our actual global conscious experiences from only this starting point. Knowing a range of conditionals of the form, *if my global conscious experience is like this, then adding this local experience will result in a global experience like that*, does not help us unless we know which of the antecedents is true.

It seems that if strong-context-dependence is true, we can account for the third platitude only on the assumption that local experiences are not (always) epistemically prior.

## 7. Taking Stock

It is time to take stock of what the discussion has shown. First and foremost: we have considered three very different models for how combination might take place, and have seen that each merits being assigned non-negligible credence at the start of enquiry. Important aspects of our concepts are captured in the platitudes we have been discussing, and each of the very different models can account for all of them. That means that the credences should remain so assigned, even after some significant conceptual investigation. Given how very different the models are, that is quite surprising.<sup>20</sup>

---

[20] As noted, Barry Dainton argues for the related but distinct conclusion that phenomenal characters of all our local experiences are unique to the precise phenomenal context in which they are found, so that a specification of any one local experience as it occurs in a global experience must make reference to *every other* local experience which contributes to the global experience. As I understand him, he argues as follows (Dainton, 2008, pp. 48, 278–9, *et passim*, the following is partly translated into my own parlance): ‘When two local experiences are “experienced together”, this has a phenomenal consequence. The only options for there to be a phenomenal consequence is for there to be an added ingredient, or for some ingredient to be modified. This consequence does not consist in some “additional ingredient”, with its own phenomenal properties, being added to the subject’s experience (on pain of the just-more-content objection; Hurley, 2002). So, the phenomenal consequence must consist in the modification of at least one constituent experience. But the “experienced together” relation is symmetrical, and holds between all parts of the global experience. So, the phenomenal consequence of local experiences constituting a global experience must be that *all* local experiences are modified by being experienced together with the other contributors.’ I remain unpersuaded by this argument because there is a third option, in addition to the two that are recognized in the second premise, namely that what accounts for the phenomenal consequence, what that consequence *con-*

It would have been easy to think that our concepts of local experiences commit us to the view illustrated by the first (or, perhaps, the second) of our models, where each local experience is a recognizable part of the global experience, the same (or very similar) each time it occurs.

But the discussion has shown that our concepts are much more permissive than this. Global experiences can share something significant enough to account for what our concepts commit us to even if they are not similar ‘on the surface’; that is, in resultant phenomenal character. As the analogies of increasing the volume of music and increasing illumination on a visual scene make vivid, they can result (partly) from the addition of a local experience, which changes a range of global experiences in a systematic way. A characterization of this scenario that springs easily to mind emphasizes dynamics: as we might say, the global experiences can share something important in virtue of resulting from a similar kind of change. But equally important is counterfactual information: to know the character of a local experience — to know what it is like to smell a rose, say — is to know how that experience would have changed a range of global experiences.<sup>21</sup>

The permissiveness of our concepts of local and global experiences is interesting. It means, among other things, that when we talk about the characters of experiences — and such talk is ubiquitous in recent philosophy, but also in everyday life — we are committed to less than we might have thought. It is important, however, to acknowledge that this comes at a cost.<sup>22</sup> Pretty clearly, it is the no-context-dependence model which captures the pre-theoretic idea of local conscious experiences most closely: there are phenomenal ‘pixels’ on our overall experiences, which one can know about, and which build up global experiences in a straightforward way.

On strong-context-dependence this picture is changed quite radically. We have found ways to account for the phenomenal character of local experiences, either as the colours which emerge in uncluttered contexts, or as the patterns of change in global experiences. But in doing so we lose some of the grip that we (thought we) had on the nature of the entities of which these are properties. I have attempted to argue that a local experience on the strong-context-dependence view saves what is worth saving from our pre-theoretic concept. But I want

---

*sists in*, is simply the existence of a global experience. Thanks to Tim Bayne for discussion.

[21] Local experiences might sometimes fail to impact a global experience at all; imagine being lightly tickled with a feather during grievous torture.

[22] Many thanks to Gabriel Rabin here.

to be up front that it does not save everything. On this conception, local experiences might be identified with the patterns of change in global experiences, they might be the dispositions to bring those changes about, or something else besides.<sup>23</sup> Either way, it is clear that we are far from our pre-theoretic starting point.

A merit of this way of developing the strong-context-dependence view is that the two other models emerge as special cases.<sup>24</sup> The no-context-dependence view is just the limiting case when it is true of *all* local experiences that it would have changed *every other* global experience in exactly the way it changes the actual one. The weak-context-dependence view is that on which it is true of all local experiences that they would have changed all global experiences in much, although not exactly, the same way.

Putting things this way is useful in assessing the burden of proof. The no-context-dependence model is very often implicitly assumed. But why should it be? It is incredible that a local experience should change global experiences in completely unsystematic ways — if it did, we would not call it an *experience*. Systematicity is clearly part of our concept. But we also need an argument to accept the other end point of the scale, where *all* local experiences always bring about exactly the same change in the character of *all* global experiences. The

---

[23] In §2 above I argued that the Combination Question coexists peacefully with a range of nearby views, including both holism and atomism about conscious experience, also on the fundamentality interpretation of that view. I was there concerned with fundamentality of *objects*. It has been suggested to me that adopting the strong-context-dependence view might commit one to atomism when that is regarded as a view about *facts* (see, for example, Fine, 2012). The suggestion is that the *fact* that a local conscious experience has the phenomenal character that it does cannot be more fundamental than the *fact* that a global conscious experience has the character that *it* does, if the strong-context-dependence view is true. First, if true, this would not be damaging: the claim is that neither holism nor atomism show the Combination Question to be misguided, uninteresting, or unimportant, not (of course) that *answering* the question a certain way will fail to impact on views elsewhere, including atomism and holism. (Its potential for impacting other views is an obvious reason to be interested in the question in the first place.) But second, I don't think the case is closed. On strong-context-dependence, we can account for phenomenal character as the tinge which arises on the pond in uncluttered contexts, or in terms of change to global experiences. But what about the entities? We could develop a view on which local experiences are the dispositions to make changes to global experiences, and it is not clear that dispositions are less fundamental features of the world than objects. If dispositions are more fundamental than objects, and if we equate the phenomenal character of local conscious experiences with the pattern of change in (actual and counterfactual) global experiences as suggested, it seems that the fact that local conscious experiences have the characters they have may well be more fundamental than the fact that global experiences have they characters *they* have, since this will amount to the claim that the fact that certain dispositions manifest in certain ways is more fundamental than the fact that certain objects are altered as dispositions manifest.

[24] Many thanks to Rachael Briggs here.

middle ground, where the change local experiences bring about in global experiences can be systematized but is not always identical stands out as more reasonable.

Another question that merits comment is how local experiences are typed, and what the consequences are for the Combination Question.<sup>25</sup> The discussion so far has presupposed that local conscious experiences are repeatable entities. We are then not talking about numerically identical experiences, since occurring at two different times suffices to prevent token-identity. So the question is: in virtue of what are two numerically distinct experiences still experiences of the same type?

A salient option for typing experiences is by phenomenal character, and a natural suggestion is that *any* difference in character yields a different type of experience. After all, it is natural to think that precisely what makes a mental state an *experience* is that it has among its identity condition its phenomenal character. Experiences just are all and only the mental states which would have been different states had they had a different phenomenal character. Or so one might well think.

Such a requirement for type-identity of experiences would immediately settle the Combination Question. It would rule out any kind of context-dependence, since a different contribution to the overall character would directly imply a different contributor state. So, one salient way of individuating local conscious experiences settles the Combination Question in favour of the no-context-dependence view.

Another natural way to individuate conscious experience is by stimulus conditions: I have talked about the experience of stubbing one's toe, for example. But on this approach the conclusion that local experiences of the same type can make different contributions on different occasions is not so surprising. So one might well wonder whether there is a way of individuating local conscious experiences on which the Combination Question remains interesting.

Suppose that we numerically individuate experiences according to what has been called 'the tripartite view' (Bayne, 2010), that is, by subject, time, and phenomenal character. It is very natural to think that two experiences can be of the same type, though they occur at different times or in different subjects. Again, many of our practices and ways of talking would be hard to understand otherwise. ('We have *all* had the experience of making a fool of ourselves in front of friends *many times*; it is at once painful and liberating'; 'I've had this experi-

[25] I am grateful to Leon Leontyev for extended email correspondence about virtually all of the issues here. I am also grateful to David Chalmers for very helpful discussion.



ence before', etc.) If token-identity is determined by subject, time, and phenomenal character, and if we can vary subject and time and still have the same type of experience, type-identity must be grounded in sameness of phenomenal character. More precisely, in *similarity* of phenomenal character: whether we require that the characters be *exactly* similar or allow some small variance ought to be open for discussion.

There is a merely verbal question which it is important to avoid (Chalmers, 2011). No one should care whether we *label* one or the other of these categories 'local experiences'. But there is also a nearby question which is *not* merely verbal. For there is a fact of the matter as to which candidate typing — that which demands exact similarity in character, that which allows slight difference, that which allows for a little more difference, and so on — plays the most significant roles in our various explanatory projects. Playing such roles is one of the most important pieces of evidence we can have for the category constituting a psychological kind — a kind which cuts the mind at its natural joints (*Phaedrus*, 265e — e.g. Plato, 2002). It seems very plausible that exact similarity in phenomenal character will not be required. For example, it is surely plausible that evidence for the very same beliefs can be constituted by perceptual experiences with slightly different phenomenal characters. By the same token, quite a lot of similarity will be required, otherwise a (slightly) different belief would be justified.

So we can and should avoid both of the ways of individuating local conscious experiences which render the Combination Question uninteresting. An important role is played by stimulus conditions: similarity in these is plausibly (at least a large part of) what singles out token experiences as candidates for being of the same type to begin with. It is an open question, at the start of enquiry, whether there is or is not something it is like to stub one's toe.<sup>26</sup>

However, stimulus conditions are not the ultimate arbiters. Sufficient similarity in phenomenal character plays this role. Whether there is or is not something it is like to stub one's toe depends on whether that stimulus elicits conscious experiences which are similar *enough*

[26] Whether there is something it is like *for Lisa* to stub her toe depends on whether a large enough proportion of *her* token experiences associated with the stimulus are similar enough in phenomenal character. Whether there is something it is like *for humans* to stub one's toe depends on whether enough of the token experiences associated with that stimulus across human subjects are similar enough in phenomenal character. If there is something it is like for humans to stub one's toe, we usually take a species-centric perspective, and say simply that there is something it is like to stub one's toe, dropping the relativization.

in character. What counts as similar enough may turn out to vary for different purposes: there may be different roles that need playing, and different important categories to play them. Or it may be that all the significant roles pick out, closely enough, the same category, so that we end up with a univocal concept of local conscious experiences.

The area of enquiry I have been trying to outline encompasses the nature of local experiences, the nature of global experience, and the nature of the relation between them. The investigation takes place within the boundaries set by the central platitudes we accept regarding all three. I have argued that for all these platitudes say, each of the three models of combination may be true. And I have argued that, for all these platitudes say, local and global experiences may turn out to be very different manners of beast. These two points are, I think, just two sides of the same coin.

### *Acknowledgments*

For discussion I am indebted to Adrian Alsmith, Tim Bayne, Sharon Berry, Rachael Briggs, Eric Brown, Einar Duenger Bøhn, David Chalmers, Alan Hájek, Mette Kristine Hansen, Uriah Kriegel, Farid Masrour, Morten Nag Opsal, Alexander Sandgren, Susanna Siegel, Sveinung Sundfør Sivertsen, Ole Martin Skilleås, Daniel Stoljar, Weng Hong Tang, Sebastian Watzl, and audiences at the Australian National University and the 2013 Unity Workshop in Bergen. Many thanks to Gabriel Rabin, Leon Leontyev, Barry Smith, and a *JCS* referee for written comments. Special thanks to John Cusbert for crucial exploratory conversations early on, out of which many of the central ideas grew.

### **References**

- Bayne, T. (2010) *The Unity of Consciousness*, Oxford: Oxford University Press.
- Bayne, T. & Chalmers, D. (2003) 'What is the unity of consciousness?', in Cleeremans, A. & Frith, C. (eds.) *The Unity of Consciousness*, Oxford: Oxford University Press.
- Bennett, D. & Hill, C. (eds.) (forthcoming) *Sensory Integration and the Unity of Consciousness*, Cambridge, MA: MIT Press.
- Chalmers, D. (2011) 'Verbal disputes', *The Philosophical Review*, **120**, pp. 515–566.
- Cleeremans, A. & Frith, C. (eds.) (2003) *The Unity of Consciousness*, Oxford: Oxford University Press.
- Dainton, B. (2008) *The Phenomenal Self*, Oxford: Oxford University Press.
- Dainton, B. (2010) 'Phenomenal holism', *Royal Institute of Philosophy Supplement*, **67**, pp. 113–139.
- Fine, K. (2012) Guide to ground, in Correia, F. & Schnieder, B. (eds.) *Metaphysical Grounding*, pp. 37–80, Cambridge: Cambridge University Press.
- Hurley, S.L. (2002) *Consciousness in Action*, Cambridge, MA: Harvard University Press.

- Jousmäki, V. & Hari, R. (1998) Parchment-skin illusion, *Current Biology*, **8**, pp. 190–191.
- Koksvik, O. (2011) *Intuition*, PhD thesis, Australian National University, [Online], www.koksvik.net.
- Koksvik, O. (forthcoming) Phenomenal contrast: A critique, *American Philosophical Quarterly*.
- Konkle, T., Wang, Q., Hayward, V. & Moore, C.I. (2009) Motion aftereffects transfer between touch and vision, *Current Biology*, **19**, pp. 745–750.
- Lee, G. (forthcoming) Experiences and their parts, in Bennett, D. & Hill, C. (eds.) *Sensory Integration and the Unity of Consciousness*, Cambridge, MA: MIT Press.
- Lewis, D. (1972) Psychophysical and theoretical identifications, *Australasian Journal of Philosophy*, **50**, pp. 249–258.
- Macpherson, F. (2011) Taxonomising the senses, *Philosophical Studies*, **153**, pp. 123–142.
- Masrour, F. (forthcoming) Unity of consciousness: Advertisement for a Leibnizian view, in Bennett, D. & Hill, C. (eds.) *Sensory Integration and the Unity of Consciousness*, Cambridge, MA: MIT Press.
- Plato (2002) *Phaedrus*, Waterfield, R. (trans.), Oxford: Oxford University Press.
- Pautz, A. (2013) Does phenomenology ground mental content?, in Kriegel, U. (ed.) *Phenomenal Intentionality*, pp. 194–234, Oxford: Oxford University Press.
- Schaffer, J. (2010) Monism: The priority of the whole, *Philosophical Review*, **119**, pp. 31–76.
- Shaw, J.R. (no date) *The Combinatorics of Conscious Experience*, unpublished manuscript.
- Schwitzgebel, E. (2008) The unreliability of naïve introspection, *Philosophical Review*, **117**, pp. 245–273.
- Smart, J.J.C. (1959) Sensations and brain processes, *The Philosophical Review*, **68**, pp. 141–156.
- Smith, B. (2012) Complexities of taste and flavour, *Nature*, **486**, S6.
- Tye, M. (2003) *Consciousness and Persons*, Cambridge, MA: MIT Press.
- Watzl, S. (2014) Attentional organization and the unity of consciousness, *Journal of Consciousness Studies*, **21** (7–8).
- Zampini, M. & Spence, C. (2004) The role of auditory cues in modulating the perceived crispness and staleness of potato chips, *Journal of Sensory Studies*, **19**, pp. 347–363.