

2

Phenomenal Unity: Mereology

In the previous chapter I introduced the term ‘phenomenal unity’ for the relation of conjoint phenomenality—the unity relation that permeates experience and provides us with a plausible conception of what it is for consciousness to be unified. Discussions of the unity of consciousness often treat phenomenal unity (or ‘co-consciousness’, as it is sometimes called) as a primitive. That is perfectly fine for some purposes, but a thorough treatment of the unity of consciousness needs to give an account of phenomenal unity. My treatment of phenomenal unity comes in two instalments. This chapter develops and defends a mereological account of phenomenal unity, according to which conscious states are phenomenally unified in virtue of the fact that they occur as the parts of a single conscious state. The following chapter examines the question of whether phenomenal unity can (also) be understood in terms of relations between the representational contents of consciousness.

2.1 The mereological model

In seeking to account for phenomenal unity it is natural to invoke the notion of subsumption (Bayne & Chalmers 2003). We might say that two conscious states are phenomenally unified when, and only when, they are co-subsumed. What it is to experience a headache and the sound of a trumpet together—what it is for these two experiences to possess a ‘conjoint phenomenal character’—is for there to be a single experience that in some way includes both the experience of the headache and that of the trumpet. Whereas treating phenomenal unity as a primitive provides us with a ‘bottom-up’ approach to the unity of consciousness, one that starts with the multiplicity in consciousness, taking subsumption as our primitive is to adopt a ‘top-down’ approach to the unity of consciousness and begin with the unity that subsumes this multiplicity.

How should we think of subsumption? It is tempting to think of it in mereological terms—that is, in terms of parts and wholes. What it is for one

experience to subsume another is for the former to contain the latter as a part. My total experiential state is a whole that includes within itself various experiential parts, such as my overall perceptual experience, my overall auditory experience, and my experience of the diesel trucks outside my window. One's overall phenomenal field is an experience that contains within itself other experiences, nestled like Russian dolls within each other. Indeed, we might venture the thought that total phenomenal states are homoeomerous: all the parts of which they are composed share their experiential nature.

Mereological language is often used in connection with the unity of consciousness. To take just a couple of examples, Lockwood says that experiences are co-conscious 'when they are parts of a complex experience' (Lockwood 1989: 88); Siewert suggests that phenomenal unity 'normally relates the constituent experiences of a single "visual" field to one another, as well as those making up a single temporal "stream" of thought and imagery' (Siewert 2001: 548); while Shoemaker holds that 'conscious states are co-conscious when they are parts of a unified state of consciousness' (Shoemaker 2003: 59). But although the mereological approach has some claim to be regarded as the standard conception of phenomenal unity it has rarely been developed in any detail. The aim of this chapter is to do precisely that.

2.2 Experience: the tripartite account

The mereological view treats phenomenal unity as a relation between token experiences—that is, between particular mental states or events.¹ Some token experiences—such as my headache experience and my auditory experience of the trumpet—are parts of a single composite experience and hence phenomenally unified with each other; other token experiences—such as my experience of the trumpet and *your* experience of it—are not phenomenally unified with each other, for there is no experience that contains both of these experiences as parts.

This account is at odds with certain conceptions of experiences, such as that which Tye develops under the 'one-experience' label. According to Tye, the only experiences that human beings have are entire streams of consciousness, where a stream of consciousness is 'a period of consciousness between one state of unconsciousness and the next' (2003: 97). The mereological model would be untenable if Tye's conception of experience were right; in fact, Tye's account *entails* that no two experiences can be phenomenally unified with each other. Tye provides

¹ I do not invest anything in the distinction between states and events, and use one term rather than the other only on grounds of stylistic convenience.

both ‘positive’ and ‘negative’ arguments for his view. His negative arguments are intended to undermine the appeal of ‘subsumptive’ approaches to phenomenal unity. I examine those arguments in §2.3. Here, I focus on the positive motivation that he provides for identifying experiences with entire streams of consciousness.²

Tye’s conception of experience is *revisionary*. It is not an analysis of our pre-theoretical notion of experience, for ordinary thought has no difficulty in taking a stream of consciousness to contain multiple experiences, both at a time and through time. This fact doesn’t mean that it should be rejected—after all, revisionary analyses can be well motivated—but it does mean that his proposal starts off on the back foot: we should identify experiences with streams of consciousness only if there are good reasons to do so.

Central to Tye’s case for the one-experience view is a supposed parallel between experiences on the one hand and clouds and statues on the other (2003: 30, 99). Clouds contain undetached collections of water molecules as proper parts, but such collections are not (typically) themselves clouds. Similarly, statues of clay may contain as proper parts undetached chunks of clay, but such chunks don’t (typically) constitute statues in their own right. Tye grants that experiences can have experiential ‘stages’, but he claims that experiential stages are no more bona fide experiences than undetached cloud parts are bona fide clouds or undetached parts of statues are bona fide statues.

I don’t find this proposal persuasive. First, a picky point: although clouds (statues) don’t typically contain clouds (statues) as proper parts, there is no principled reason to deny that they *could*. One can build a statue that contains other statues as parts. But Tye is certainly right to point out that we don’t typically regard undetached regions of water molecules contained within clouds as clouds in their own right, nor do we regard undetached regions of clay contained within statues as statues in their own right. So why should we regard arbitrary components of a stream of consciousness as experiences in their own right?

Well, I’m not sure that we *should* regard every undetached component of a stream of consciousness as an experience in its own right. Consider the stream of consciousness that began when you woke up this morning and that will conclude when you fall into a dreamless sleep tonight. The component of this stream that occupied the first waking hour of your day is not itself an experience—at least not in the sense that I am considering here. I don’t deny that there is *a* use of the term ‘experience’ that can be used to pick out arbitrarily large stretches of consciousness. (‘Tell me about your time in the Gobi desert.’ ‘It was

² For further discussion of Tye’s position see Bayne (2005) and Dainton (2004).

great—a really fantastic experience.’) In this very broad sense of ‘experience’ the chunk of my stream of consciousness that occupied the first waking hour of today might indeed qualify as an experience in its own right. But this is not the notion of experience that I am after here. Experiences, as I am interested in them here, are states that can be enjoyed ‘all at once’.

Nonetheless, although I would resist the claim that *every* part of a stream of experience is an experience in its own right, it does seem to me that *many* parts of a stream of experience qualify as experiences in their own right. My current conscious state is an experience in its own right, as are the various fine-grained conscious states—the pain in my left leg; my olfactory experience of the coffee; my auditory experience of the dog barking in the alley—that are contained within it.

We should not be misled by Tye’s discussion of clouds and statues, for there are important differences between clouds and statues on the one hand and experiences on the other. For one thing, clouds and statues are physical objects, whereas experiences are events. Even if it is generally true that the proper parts of an *object* of kind K do not constitute objects of kind K in their own right (and there are exceptions to this generalization—consider computers), *events* often contain as proper parts events of the same kind. Arguments contain as proper parts other arguments, battles contain as proper parts other battles, traffic jams contain as proper parts other traffic jams, and stories can contain as proper parts other stories. (Consider *Hamlet*.)

Furthermore, there is usually little need to distinguish the parts of a cloud from each other or the parts of a statue from each other, but there is indeed some point in distinguishing the parts of a stream of consciousness from each other. For one thing, we need some way of referring to the distinct instances of a particular experiential type that occur within a single stream of consciousness. Consider again the stream of consciousness that began when you woke up this morning and will conclude when you fall into a dreamless sleep tonight. Now, suppose that this stream contains one instance of headache phenomenology at its beginning and another towards its end. It is natural to say that this stream of consciousness contains two headache experiences. Indeed, it becomes very difficult to describe the internal structure of a stream of consciousness if we insist on identifying experiences with entire streams of consciousness. Tye, of course, could say that such a stream of consciousness contains two stages that include headache phenomenology, but at this point his ‘experiences stages’ would appear to be experiences in all but name. I suppose that one could insist on distinguishing experiences from experiences stages, but I don’t see anything to be gained from doing so. All in all, I see no reason to follow Tye in reserving

the term ‘experience’ for entire streams of consciousness. In fact, on the conception of experiences that I will embrace, the typical stream of consciousness does *not* constitute an experience in its own right.

Tye’s one-experience conception of consciousness can be set to one side, but it does raise the important question of how we should think of token experiences. I am not convinced that there is any single way in which experiences should be individuated. Counting experiences is arguably more like counting the number of objects in a room or the number of events that took place during a meeting than it is like counting the number of beans in a dish: one has some idea of how to go about one’s business, but the idea that there is only one way in which to proceed is somewhat farcical. The notion of a token experience is elastic, and different approaches to the individuation of experiences might be appropriate in different contexts.

That said, I do think that there is one conception of experiences that is particularly well-suited for addressing the questions raised by the unity of consciousness: the *tripartite conception*. According to the tripartite conception of experience, experiences are to be individuated in terms of subjects of experience, times, and phenomenal properties.³ In other words, token experiences must differ from each other in terms of whose experiences they are, when they occur, or the kinds of phenomenal properties that they involve. Although I have introduced the tripartite analysis by reference to phenomenal properties, we can also think of it in terms of phenomenal events, for events can be understood in terms of the instantiation of properties (Kim 1976).

There is a natural fit between the tripartite conception of experience and the mereological conception of phenomenal unity, for within the stream of consciousness we can identify both more and less complex experiences. An experience produced by tasting a strawberry will involve a range of distinct phenomenal properties, such as ‘tanginess’, ‘sweetness’, and—dare one say it—‘strawberryness’. We can think of each of these phenomenal properties as involving distinct experiences that are parts of the more complex experience of tasting the strawberry. That experience, in turn, will be a part of even more complex experiences. Some of these more complex experiences will be modality-specific, others—such as the experience that corresponds to one’s overall perceptual phenomenology—might include content drawn from multiple modalities. Indeed, one’s overall phenomenal field—what it is like to be you right now—is, in

³ The notion of temporal location that I am working with here allows both temporal points and brief periods of time. I suspect that there are no instantaneous experiences, but I see no reason to assume that here. By a brief period I mean something that is no longer than a specious present (however long that is).

my view, a very complex phenomenal event that contains within it the rest of one's experiential states.

According to this conception of experiences, not everything that has experiences as parts is itself an experience. The typical stream of consciousness—a period of consciousness between one state of unconsciousness and the next—is not best thought of as a phenomenal event, for there is no single phenomenal property that corresponds to a typical stream of consciousness. There is something it is like to enjoy a typical stream of consciousness, but this 'what it's likeness' is spread out—distributed across a number of distinct conscious states. It lacks the kind of unity that the phenomenal field possesses. (The only stream of consciousness that might qualify as a phenomenal event would be a very short stream—a stream whose duration was no longer than a single specious present.)

It might be objected that the tripartite account of experiences has counter-intuitive consequences—that it individuates experiences in ways that are significantly at odds with our ordinary ways of counting experiences. Of course, one might wonder just how much weight our pre-theoretical practices ought to be afforded here, but waiving that concern let us examine just what kind of friction there might be between the tripartite approach and 'common sense'.

Consider first colour experience. Suppose that you are experiencing a certain shade of blue. According to the tripartite account, there will be an experience corresponding to the event of your instantiating this phenomenal property. But, intuitively, one could have multiple experiences of blue at a single point in time, for one can see multiple objects and regions of space as being blue. Here, the tripartite account appears to entail that one would have only a single experience whereas there is some pre-theoretical force to the thought that one actually has multiple experiences of blue.

There are various ways in which we might respond to this objection, but perhaps the best response is to restrict the tripartite account to maximally specific or fine-grained phenomenal properties, where a phenomenal property is maximally specific if it has no determinates (Bayne & Chalmers 2003). The phenomenal property of blue is a determinable that has as determinates the phenomenal property of blue occurring in a certain location of space, but arguably this phenomenal property has no determinates. And there is an intuitive sense in which one couldn't have multiple instantiations of *that* phenomenal property. Once restricted to fine-grained, maximally determinate phenomenal properties the tripartite account faces no objection from this quarter.

Another objection to the tripartite account derives from considerations concerning the 'common sensibles'—that is, properties that can be detected via more than one modality. Consider a fundamental common sensible:

motion. Suppose that you are watching an ant crawl across your skin. You are aware of the ant's movement in two ways: via vision and via touch. Yet—the objection runs—although there are intuitively two experiences of motion here, the tripartite account entails that there is only a single experience, for there is only one phenomenal property in question: a representation of motion. It looks as though we will need to appeal to some further element over and above the three components of the tripartite model in order to account for the contrast between visual experiences of motion and tactile experiences of motion.

One line of response to this objection would be to argue that the visual experience of motion involves one phenomenal property while the tactile experience of motion involves another. Generalizing this response, we might say that when different senses represent the same property they will often do so via distinct ways—that phenomenal properties are more fine-grained than the worldly properties that they represent. This 'Fregean' conception of phenomenal properties will allow that experiences can be grouped into modality-specific clusters without putting any pressure on the tripartite account.⁴

Although the approach just outlined will have its advocates, I prefer a rather different response to it. I think the objection is best met by biting the bullet: in the situation outlined one *would* have only a single experience of the ant's motion. I'm inclined to think of phenomenal properties in 'Russellian' terms—that is, in terms of 'external world' properties such as motion itself.⁵ On this view, what it is like to see the ant as moving at such-and-such a speed is identical to what it's like to feel it as moving at such-and-such a speed, for the two cases involve the representation of one and the same property—namely, movement at such-and-such a speed. On this view, the experience of the ant's motion is 'amodal' when considered in and of itself. Situations in which one is both visually and tactually aware of the ant's motion are not best described as situations in which one has two experiences of the ant's motion, but as situations in which one has a single experience of the ant's motion that is supported by two streams of perceptual processing (each of which is redundant given the presence of the other).

Of course, there *are* differences between seeing an ant move and feeling it move. For one thing, sight and touch will rarely have the same degree of resolution. Moreover, even if these two senses were to represent exactly the same speed of the ant, they would do so in the context of other modality-specific representations. For example, the visual representation of the ant's motion will occur in the context of a visual representation of its colour, whereas

⁴ See e.g. Chalmers (2004), Thompson (2009), Kulvicki (2007).

⁵ See e.g. Dretske (1995); Harman (1990); Tye (2003).

the tactile representation of the ant's motion will occur in the context of a representation of one's own body. In this way the Russellian can account for our ability to say whether an experience of an ant's motion is visual, tactile, or both visual and tactile. The upshot of the foregoing is that the tension between the tripartite analysis and our ordinary ways of counting experiences is more apparent than real.

Before concluding this section, let me contrast the tripartite approach with another approach to the individuation of experiences, an approach that is perhaps its natural rival. The approach in question appeals to the physical-functional basis of experience in order to segment the stream of consciousness into parts. To invoke a term that has some currency in the current literature, we might try to individuate experiences in terms of their *vehicles*.

Although there may be perspectives which demand a vehicular approach to the individuation of experience, I think we have reason to prefer the tripartite approach when it comes to questions concerned with the unity of consciousness. Experiential states are states of organisms—they are not states of hemispheres or parts of brains. There is, of course, a sub-personal account to be given of why a particular organism might enjoy experiential states at any one point in time, and of why it might enjoy the particular experiential states that it does, but we ought to be wary of the suggestion that the sub-personal features of consciousness might enter into their identity conditions. The unity of consciousness is an *experiential* aspect of consciousness, and our approach to it must take this fact seriously.

A second concern with vehicular approaches to consciousness is somewhat pragmatic in nature. Although there is much talk of the vehicular nature of consciousness, there are rather few worked-out accounts of just what the vehicles of experience are supposed to be or how one might go about pinning them down. The problem here is not (just) that we know so little about the neuro-functional basis of consciousness—although that certainly doesn't help—but that the very notion of a vehicle of consciousness remains obscure. On some accounts the vehicles of consciousness are identical to conscious states, and phenomenal properties are simply entities that those states 'carry'. On other accounts the vehicles of consciousness are not identical to experiences but they do constitute them or form their supervenience base. On still other accounts, vehicles are neuro-functional states that underwrite a certain kind of explanation of consciousness. Even within the explanatory conception of vehicles there are various positions that might be adopted, depending on the kind of explanation that one is looking for (Hurley 2010). Given this confusing welter of conceptions, it is very difficult to say just what a vehicular approach to the

individuation of experiences might look like. Would it individuate experiences in fine-grained terms, say, according to perceptual qualities such as colour, shape, motion, timbre, and so on, or would it individuate experiences in coarse-grained terms, according to (say) perceptual objects, particular modalities, or indeed entire phenomenal fields? Not only do we lack good answers to these questions, it is rather unclear how we ought to go about determining the answers to them.

Certain readers might be inclined to respond that cognitive neuroscience has already identified the vehicles of consciousness. Don't we know—some may say—that V4 functions as the vehicle for the visual experience of shape (Pasupathy & Connor 1999, 2001); V4/V8 functions as the vehicle for the experience of colour (Bartels & Zeki 2000; Hadjikhani et al. 1998); and V4 (MT/MST) functions as the vehicle for the experience of motion (Tootell, Reppas, Dale, et al. 1995; Tootell, Reppas, Kwong, et al., 1995)? But whether indeed this is something that is known depends very much on just what vehicles are supposed to be. These regions may indeed play a particularly important role in the generation of states with the appropriate contents—for example, activation in V4 might explain why the creature in question is enjoying experiences of shape rather than some other type of experience—but we should not *identify* such experiences with activity in V4 (see also §10.5). Indeed, it is implausible to suppose that V4 activity is even sufficient for visual experiences (as) of shape. A slice of V4 that has been placed in a test tube won't generate visual experiences no matter how much current is run through it. I suspect that even with a completed cognitive neuroscience of consciousness, we will still be left with competing ways of carving up the stream of consciousness into distinct experiential vehicles.

Let us recap. The business of counting experiences is a messy one, and there is more than one respectable way of going about it. That being said, I suggested that experiences should be thought of in tripartite terms: an experience is to be understood in terms of the instantiation of a phenomenal property by a subject at a time. We can think of these instantiations as phenomenal events. And in light of this, phenomenal unity can be understood in terms of mereological relations between phenomenal events. At any one point in time one's stream of consciousness takes the form of a single highly complex phenomenal event that subsumes a number of less complex phenomenal events. It is the fact that these less complex events are proper parts of more complex events that accounts for their unity. The mereological relations between phenomenal events might be reflected in mereological relations between their vehicles but they need not be.

2.3 Objections and replies

Let us consider now some objections to the mereological account. I begin with Tye's objections to the view, before turning to criticisms from Brook and Raymont on the one hand and from Searle on the other.

The self-undermining objection

A first objection to the mereological account holds that it is in some way self-undermining:

Consider a maximal phenomenal state (e_m); that is, a phenomenal state that is not subsumed by any other phenomenal state. e_m includes as two of its elements a visual experience (e_1) and an auditory experience (e_2). Not only are these two experiences unified with each other, but each is unified with e_m . Now, if the unity of each of these two experiences requires that there be a unifying experience that subsumes them, then the unity of e_1 (e_2) and e_m seems to demand that the subject have a further experience, 'bigger' than e_m , that subsumes both e_m and e_1 (e_2). But we stipulated that e_m was a maximal phenomenal state, a state not subsumed by any other phenomenal state. So if phenomenal unity is a relation between experiences, then the notion of a maximal phenomenal state is incoherent. But the notion of a maximal phenomenal state clearly *is* coherent, so phenomenal unity cannot be a relation between experiences. (Tye 2003: 22)⁶

The notion of a maximal (or 'total') phenomenal state is indeed coherent, but why exactly must the proponent of the mereological account deny this? The crucial move in the argument is the claim that, on the subsumptive view, the unity of e_1 and e_m requires that there be an experience that is 'bigger' (that is, contains more phenomenal content) than e_m itself. But we are under no obligation to accept this claim, for parthood is a reflexive relation—any event is an (improper) part of itself. The state that unifies e_m and e_1 need be none other than e_m itself.

As an analogy, think of what it is for two actions to be 'unified'. We might regard writing a book as an act that subsumes the act of writing the book's first sentence. What it is for these two actions to be unified is for the latter to be included as a component or proper part of the former. Nonetheless, we can still think of writing a book as a maximal action that is not—or at least need not be—subsumed by any more complex action. Similarly, we can think of a total phenomenal state as a state that subsumes each of the phenomenal states that the subject has at the time in question without being subsumed by any state other than itself.

⁶ I have substituted Tye's nomenclature for my own.

The 'phenomenal bloat' objection

A second objection holds that the mereological account is faced with the problem of what we might call *phenomenal bloat*. Various versions of this objection have been circulating in the literature (see Hurley 1998; Bayne 2001), but I will focus on Tye's version of it.

- (1) Suppose that phenomenal unity is a relation between experiences e_1 – e_5 (assumption for *reductio ad absurdum*).
- (2) This unity relation (R_1) between experiences must itself be experienced, for if there were no experience of the unifying relation, then there would be nothing it is like for e_1 – e_5 to be unified. (Alternatively, there would be no phenomenal difference between a situation in which e_1 – e_5 were phenomenally unified and a situation in which they are not unified.)
- (3) If R is itself experienced, it must have its own phenomenology.
- (4) If R has its own phenomenology, its phenomenology must be unified with that of e_1 – e_5 .
- (5) In order to account for the fact that R_1 is unified with e_1 – e_5 we need to posit another unity relation (R_2).
- (6) But of course R_2 must itself be experienced, for if there were no experience of the unifying relation, then there would be nothing it is like for e_1 – e_5 to be unified with R_1 .
- (7) But now we have embarked on a vicious infinite regress.
- (8) So (1) must be false: phenomenal unity is not a relation between experiences. (Tye 2003: 22)

The first thing to say about this argument is that if it is any good it seems to be too good, for variants of it threaten to undermine *any* account of phenomenal unity. If phenomenal unity is not a relation between experiences, it looks like it must be a relation between the contents of experience, as Tye has claimed (see §3.1). And what could make it the case that simultaneously experienced perceptual qualities—the loudness of a sound, the smoothness of a surface, and the sweetness of a taste—enter into the same phenomenal content? Arguably, these qualities could enter into a single content only if there were an experiential difference between experiencing the loudness of a sound, the smoothness of a surface, and the sweetness of a taste together as opposed to experiencing these properties separately. But—the objection continues—this surely entails that the subject must be conscious *that* they are conscious of the sound, the surface, and the taste, and if that's right then we face the task of explaining how *that* experiential content is unified with the rest of the subject's phenomenology. We appear to have embarked on an apparently vicious regress, not unlike the one outlined above. It is no accident that Hurley's (1998) version of the phenomenal bloat objection—which she dubbed the 'just more content' objection—targeted content-based accounts of the unity of consciousness.

Leaving aside the question of whether content-based conceptions of phenomenal unity might face their own problem of phenomenal bloat, how is the mereological account to respond to it?

One line of response to the argument is to reject premiss (5). We might suppose that although phenomenal unity is an experience it is a peculiar sort of experience: it is a *self-binding* experience. Unique amongst experiences, perhaps phenomenal unity binds experiences together to form unified phenomenal wholes without itself needing to be so bound. However, I am not much attracted to this objection, for why should some experiences be self-binding and others not?

Another line of response would be to reject (3). One might hold that although relations of phenomenal unity are experienced, the experience of them does not possess a distinctive phenomenal character of its own. Christopher Hill may have had something like this view in mind when he suggested that there is a kind of ‘pure’ or ‘ghostly’ element in consciousness, an element that ‘has no distinguishing characteristics other than its ability to unite sensations’ (Hill 1991: 239). However, Hill came to reject this view, and I think he was right to do so. For one thing, phenomenal unity lacks the kind of positive character that would allow it to be singled out via introspection. We seem unable to attend to phenomenal unity as such. One can attend to phenomenally unified experiences, and one can attend to the conjoint experience that subsumes them, but one cannot attend to relations of phenomenal unity or subsumption themselves. Here is another way to appreciate the difference between phenomenal unity and experiences. Unlike experiences, relations of phenomenal unity cannot themselves be phenomenally unified. But, since any experience is the sort of thing that could be phenomenally unified, it follows that relations of phenomenal unity are not experiences. (Judgements about relations of phenomenal unity might have their own phenomenology, but that’s a different matter.)

In my view, the argument’s flaw lies with premiss (2)—the claim that in order to make a phenomenal difference phenomenal unity must itself be experienced. Tye provides the following argument for (2): ‘If there were no experience of the unifying relation, then there would be nothing it is like to have the sense specific experiences unified’ (2003: 22). I think this inference should be resisted. Why could there not be something it is like to have a set of unified experiences, without that ‘what it’s like’ subsuming or involving an experience *of* the unity relation that binds the experiences in question together? Phenomenal unity is a phenomenal relation in the sense that it makes a phenomenal difference, but not in the sense that it has its own phenomenal character that makes an *additional* contribution to what it is like to be the subject in question. We can think of this in terms of the different

ways of undergoing experiences e_1 – e_5 . In principle, one can have these experiences separately or one can have them together, as parts of a subsuming experience. Unity then is not an *object* of experience but a *manner* of experiencing.⁷ The lesson to be learnt from the phenomenal bloat objection is that phenomenal relations can make a phenomenal difference without themselves being either experiences or experienced.

The transparency objection

Rather than attack the mereological account head-on, one might attempt to undermine its epistemic basis. In fact, Tye has done precisely that by appealing to the so-called transparency (or diaphanousness) of consciousness. Introspection, says Tye, gives one access only to the contents of experience. And this fact, says Tye, undercuts any account of phenomenal unity that treats it as a relation between token experiences:

Visual experiences are transparent to their subjects. We are not introspectively aware of our visual experiences any more than we are perceptually aware of transparent sheets of glass. If we try to focus on our experiences, we ‘see’ right through them to the world outside . . . If we are not aware of our experiences via introspection, we are not aware of them as unified. The unity relation is not given to us introspectively as a relation connecting experiences. Why, then, suppose that there is such a relation at all? (Tye 2003: 24 f.)⁸

The so-called transparency of consciousness is a rather complex matter, and appeals to it must be handled with care.⁹ Consider first the bold claim that we are not aware of *any* of our conscious states via introspection but only of their contents. If the notion of a ‘conscious state’ is understood to include thoughts then this claim is surely implausible, for there is an introspectible difference between consciously judging that it is sunny and consciously hoping that it is sunny. And if introspection can distinguish thoughts that differ only with respect to their mode or attitude, then it must have access to more than the contents of thoughts. Of course, Tye’s claim here is only that *visual* experiences are transparent to their subjects, but once we have allowed that introspection can distinguish between thoughts with the same content we have some reason to wonder whether it mightn’t also allow us to distinguish between perceptual experiences with the

⁷ I am indebted to Ian Phillips here.

⁸ Tye’s one-experience view of course entails that there are no purely visual experiences, so presumably he means ‘visual experiential stages’ or some-such.

⁹ The contemporary discussion of transparency begins with Harman (1990). For critical discussion see Kind (2003), Stoljar (2004) and (2007), and Thompson (2008).

same content—indeed, whether we can't even introspectively distinguish between *visual* experiences with exactly the same content. Contrast visual *perception* (as) of an apple with visual *imagery* (as) of an apple. One might argue that it is possible to distinguish these two states from each other on the basis of introspection even though they have exactly the same contents. Advocates of transparency do have responses to this objection—for example, they can argue that there must be subtle content-based differences between the perceptual state and the imagery state if indeed they are introspectively distinguishable—but the plausibility of such responses is something of an open question.

Despite these reservations, I do think that there is something to claims of transparency. We certainly do not have introspective access to the sub-personal basis of conscious states—their nature as neuro-functional states. However, whether this implies that we are not introspectively aware of our experiences depends on just how we conceive of experiences. Suppose, as I have suggested, that we identify experiences with the instantiation of phenomenal properties. On this conception of experiences, we *are* aware of our experiences as such in being aware of which phenomenal properties we instantiate. I am currently aware that I instantiate the phenomenal property distinctive of tasting coffee, and in being aware of this fact I am also aware of the corresponding experience. On the tripartite conception of experience, the contrast between 'access to experiences' and 'access to the contents of experience' is a false one, for in having access to the contents of experience one also has access to experiences themselves.

Of course, advocates of transparency will be likely to reject this conception of experiences in favour of a vehicular conception. On their view, experiences themselves are (say) neural events that 'carry' contents (whatever that means). From this perspective, introspective access to experience itself would require that one be aware of neural events as such, and it is clear that introspection involves no such awareness. But even from within this framework we have the resources to motivate the view that phenomenal unity is a relation between experiences.

Distinguish between introspective access to a *state* and introspective access to a *fact*. The vehicular conception of experience may imply that we lack introspective access to our experiential states, but it must surely allow that we have introspective access to certain facts about our experiences. For example, one might have introspective access to the fact that one is conscious, to the fact that one is visually conscious, and to the fact that one has a visual experience (as) of there being an apple in front of one. Furthermore, some of the facts that introspection provides access to concern the relations between one's experiences. The reason for this is that one has introspective access to facts about the

contents of one's consciousness, and such facts constrain the relations between one's experiences themselves. One doesn't merely have (say) an experience of an apple, a trumpet, and an itch; instead one experiences the apple, the trumpet, and the itch 'together' within a single phenomenal field. In order for these contents to be unified the experiences that underlie them must also be unified: no unity in content (or phenomenal character) without unity between the experiences that carry those contents. In short, we can have introspective reasons for thinking that there is a unity relation connecting experiences even if that relation is not directly 'given to us' in introspection.

The 'Jamesian' objection

Tye is not the only critic of the mereological approach to phenomenal unity. Joining him are Brook and Raymont, who argue that the account fails for reasons that can be found in the work of William James (and which in fact have their roots in Kant). Merely putting conscious experiences together, Brook and Raymont claim, will fail to generate a unified state of consciousness, for a 'mere combination of experiences is not the experience of a combination' (Raymont & Brook 2009: 575). Here is how James put the point:

Take a hundred [feelings], shuffle them and pack them as close together as you can (whatever that may mean); still each remains the same feeling as it always was, shut in its own skin, windowless, ignorant of what the other feelings are and mean . . . Take a sentence of a dozen words, take twelve men, and to each one word. Then stand the men in a row or jam them in a bunch, and let each one think of his word as intently as he will; nowhere will there be a consciousness of the whole sentence. (James 1890/1950: 160)

We can—indeed, we *should*—accept that the mere combination of experiences is not identical to the experience of combination. Simply packing together experiences of a number of words—'jamming them together in a bunch'—does not suffice to generate an experience of the sentence as a whole. But the mereological account does not claim that it does. The mereological account is intended only as an account of what it is for experiences to be phenomenally unified, it is *not* intended as an account of what it is for consciousness to be representationally integrated. It is not an account of what it is for experiences of (say) the parts of a face to give rise to an experience of the face to which those parts belong, for the experience of individual notes to give rise to an experience of the melody those notes realize, or for the experience of individual words to give rise to an experience of the sentence that those words compose. We obviously *do* need an account of how the experience of parts generates an experience of the whole of which those parts are parts, and it is plausible to suppose that phenomenal unity will play an important role in such an account, but we should not conflate representational

unity with phenomenal unity. Experiencing the parts of an object (scene, event) might be *necessary* in order to experience the object (scene, event) as such, but it typically isn't *sufficient* for it, and the two forms of experience should certainly not be identified.

One of the upshots of the foregoing is that we must distinguish between two kinds of binding problems (Revonsuo 1999). The binding problem is typically thought of in representational terms—that is, as the problem of ensuring that perceptual features are integrated into unified percepts of objects. Consider vision. Given that we typically experience multiple visual objects at a time, the brain needs to work out which perceptual features (colour, shape, texture, etc.) belong to which objects. Suppose that one is looking at a red berry nestled in green leaves. It is important that one represents the berry as red and the leaves as green rather than vice versa. This binding problem must be distinguished from the phenomenal binding problem—the problem of ensuring that each of the phenomenal features that one enjoys at a certain point in time occurs within a single conscious state. Even when features are not bound together as the features of a single object (as in visual agnosia), they are nonetheless experienced together within a single conscious state. James was surely right to note that merely sticking features together in a single representation—‘jamming them as close together as one can’—does not ensure that they will be experienced as the features of a single object, but it arguably does ensure that they will enjoy a conjoint phenomenal character. Here, as elsewhere, it is vital that we distinguish phenomenal unity from representational unity.

The building block objection

A final objection to the mereological view comes from Searle. Searle contrasts mereological conceptions of consciousness—which he dubs ‘building block models’—with unified field models, according to which the multiplicity to be found within consciousness ought to be understood in terms of modifications to an overall field of consciousness:

The urge to think of consciousness . . . as made up of smaller building blocks is overwhelming. But I think it may be wrong for consciousness . . . Indeed, maybe it is wrong to think of consciousness as made up of parts at all . . . Instead of thinking of my current state of consciousness as made up of the various bits—the perception of the computer screen, the sound of the brook outside, the shadows cast by the evening sun falling on the wall—we should think of all these as modifications, forms that the underlying basal conscious field takes after my peripheral nerve endings have been assaulted by the various external stimuli. (2000: 575)

I share Searle's suspicion of building block approaches to consciousness. I also share his attraction to the thought that the multiplicity in consciousness might involve the modification of an underlying basal field (see Chapter 10). But—unlike Searle—I want to distinguish the mereological conception of consciousness from the building block conception. Searle's 'building block' terminology suggests a particular conception of the relationship between the overall phenomenal field and the parts of consciousness—namely, a view on which those parts are autonomous units of consciousness, states whose status as conscious states is independent of their location in the particular phenomenal field in which they occur. No such assumption is implicit in the mereological account of phenomenal unity. To say that the phenomenal field contains experiential parts is to make no claim whatsoever about the relationship between those parts and the whole of which they are parts. We could take the experiential parts of a phenomenal field as more fundamental than the phenomenal field itself, but we could also hold that there is a sense in which the overall phenomenal field is more fundamental than its parts. Indeed, my own view—for which I argue at length in Chapter 10—is that the total phenomenal state is prior to and more fundamental than the experiential parts that it subsumes. But that claim is independent of the mereological account, which is strictly neutral on the question of whether in seeking to understand the structure of consciousness we should assign explanatory priority to the parts (Searle's 'building block' model) or to the whole (his 'unified field' model).

This completes my defence of the mereological account of phenomenal unity. I turn now to consider one of the central questions that it raises.

2.4 Partial unity

Phenomenal unity is clearly both reflexive (every state is unified with itself) and symmetrical (if e_1 is phenomenally unified with e_2 then e_2 must be phenomenally unified with e_1). But is it also transitive? If e_1 and e_2 are each phenomenally unified with e_3 does it follow that they must also be unified with each other? The mereological account as such leaves this question open, but it is a question that a full analysis of the unity of consciousness must address.

Arguably phenomenal unity need not be transitive when the states involved are not simultaneous. Consider what it's like to hear three successive notes (*Do*, *Re*, and *Me*) where these notes are played such that one's experiences of *Do* and *Re* occur together within a single specious present, one's experiences

of *Re* and *Me* also occur within a single specious present, but one's experiences of *Do* and *Me* do *not* occur within a single specious present. Such episodes appear to involve a failure of transitivity in phenomenal unity, for one's experiences of *Do* and *Me* both appear to be unified with an experience of *Re* but not with each other (Dainton 2006). But it is far less plausible to suppose that transitivity can fail for sets of *simultaneous* experiences. In fact, it is tempting to suppose that the phenomenal field *cannot* fragment in the way that a failure of transitivity would require. In other words, it is tempting to suppose that for any three simultaneous experiences, e_1 , e_2 and e_3 , if both e_1 and e_2 are phenomenally unified with e_3 then they must also be unified with each other. Let us call the assumption that phenomenal unity is transitive with respect to simultaneous states *the transitivity thesis*.

Although the transitivity thesis has a great deal of intuitive plausibility, it has not gone unquestioned. In *Mind, Brain and the Quantum*, Michael Lockwood suggested that it might be possible for a subject's consciousness to fragment into a single partially unified stream of experience in which transitivity fails. (In fact, Lockwood suggested that the partial unity model of consciousness might even capture the structure of consciousness in certain split-brain patients (see §9.4)). Lockwood's proposal has met with a mixed response. Tye claims that there is no difficulty with partial unity, although it is far from clear that his conception of partial unity is the same as Lockwood's (as we shall see). Dainton allows that transitivity might fail in the context of successive experiences, but he denies that such failures are possible when it comes to simultaneous experiences (Dainton 2006). Hurley neither accepts nor rejects the transitivity thesis, but she does provide an account of the split-brain syndrome that is designed to remove the attractions of the partial unity model and in so doing implies that we should be at least sceptical of partial unity. Indeed, even Lockwood himself came to harbour reservations about the possibility of partial unity (Lockwood 1994).

The possibility of partial unity is a matter of no little importance. As we shall see in later chapters, apparent breakdowns in the unity of consciousness are rarely clean, and there is some temptation to think that if and when the unity of consciousness is lost it is lost because consciousness takes the form of a single, partially unified stream of consciousness rather than two streams of consciousness, each of which is fully internally unified. So, is partial unity possible?

Let us begin by considering in more detail just what it would take to have a case of partial unity. Tye offers the following as an example of such a scenario:

S has two multimodal experiences, e_1 and e_2 . e_1 represents the pinprick in *S*'s neck, his left arm, fingers and the surface his left forefinger is touching. e_2 represents the pinprick in *S*'s neck, his right arm, fingers, and the surface the right forefinger is touching. The pricking is phenomenally unified with redness by entering into the same phenomenal

content—the phenomenal content of e_1 . The pricking is phenomenally unified with greenness in like manner, but this time the common content is the phenomenal content of e_2 . Since S has no experience whose phenomenal content has entering into it both greenness and redness, the two colours are not phenomenally unified. (2003: 130 f.)¹⁰

Strictly speaking, Tye's one-experience view requires that e_1 and e_2 must be experience stages rather than experiences, for neither e_1 nor e_2 is an entire stream of consciousness, but let us waive that concern here. The central question is whether this scenario does indeed involve a failure of transitivity. In my view it does not. Lockwood introduced the notion of non-transitivity in order to describe a split within the structure of consciousness itself, but—as befits his conception of phenomenal unity (see §3.1)—Tye conceives of partial unity in terms of relations between the *contents* of consciousness. Tye denies that e_1 and e_2 share a common experiential part. Instead, on his view they 'overlap' only in the sense that both experiences have a common content—namely pricking. But this is a very different conception of partial unity from that which I, following Lockwood, am considering. Partial unity, as I am thinking of it here, would occur when—and only when—the experience of pricking that occurs in e_1 is *numerically identical* to that which occurs in e_2 . On this view, e_1 and e_2 must share an experiential part—namely, an experience of pricking.

Having clarified the notion of partial unity let us return to the question of whether it is possible. In asking about the coherence of partial unity it is vital that we take into account the fact that we are addressing relations between *conscious* states. It is not difficult to make sense of the idea that there might be a kind of 'partial unity' between unconscious mental states. Consider beliefs, understood as dispositional states that play a distinctive functional role. It is clear that partial unity can obtain between beliefs thus understood. We can easily conceive of a scenario in which beliefs m_1 and m_2 are both disposed to exert control over thought and behaviour in conjunction with m_3 but not in conjunction with each other. Indeed, beliefs m_1 and m_2 might be at odds with each other, in the sense that the activation of the one belief tends to inhibit the activation of the other. (Something like this might occur in the context of self-deception.) But conscious states are not dispositional states, and the fact that we can make sense of 'partial unity' when considering dispositional states has little bearing on whether *phenomenal* states can be partially unified.

One might argue that partial unity is incoherent on the grounds that there is no consistent assignment of partially unified states to subjects of experience. Suppose that e_1 and e_2 are both unified with e_3 but not with each other. One might reason as follows. Both e_1 and e_3 must be assigned to the same subject for

¹⁰ Here too I have substituted Tye's nomenclature for my own.

they are phenomenally unified, and as phenomenally unified states they are subsumed by a single experiential state. Precisely the same considerations suggest that e_2 and e_3 must also belong to the same subject. But, so the argument goes, e_1 and e_2 *cannot* belong to the same subject of experience, for they are not phenomenally unified with each other. In short, the objection is that there is no consistent way of assigning partially unified experiences to subjects of experience.

The key move in the argument is obviously the final one: the assumption that simultaneous experiences that belong to the same subject must be phenomenally unified. There is some plausibility to this claim—in fact, it follows from the unity thesis—but we cannot assume that it is true at this stage of proceedings. The unity thesis might be true but it is not a conceptual truth, at least not if we are thinking of subjects of experience in biological terms. In fact, the unity thesis is unlikely to be a conceptual truth even on certain *psychological* conceptions of the subject of experience. For example, views that identify subjects with networks of functionally defined mental states appear to allow that there might be occasions on which a single subject has experiences that are not phenomenally unified with each other (see e.g. Shoemaker 1996). So, although there is something to the thought that a single subject could not have disunified experiences, let us assume that there is at least a viable notion of subjecthood according to which simultaneous but disunified experiences can be assigned to the same subject.

The most tempting reason to reject partial unity is simply that it is difficult if not outright impossible to project oneself into the perspective of a partially unified subject (Lockwood 1994: 95). We can convert this thought into the following argument:

- (1) If partial unity were possible then there would be something distinctive it is like to be a partially unified subject—there would be such a thing as a partially unified phenomenal perspective.
 - (2) We are unable to project ourselves into a partially unified phenomenal perspective.
 - (3) If there were such a thing as a partially unified phenomenal perspective then we should be able to project ourselves into it.
- (C) Thus, partial unity is impossible.

Let us call this *the projectability argument*. On the face of things the first premise of the argument seems secure. Although there are those who might reject it—Paul Snowdon describes the assumption that for each subject of experience there is

a way that it's like to be them as 'hazy' (1995: 78)—I see no *prima facie* reason to doubt it. But in a very provocative discussion of the projectability argument Hurley suggests that premise (1) might be rather more problematic than it appears.¹¹

Hurley argues that it is a mistake to conceive of the difference between full unity and partial unity in subjective terms.¹² Her objection centres on the claim that the contrast between partial unity and full unity is a matter of the relations between token experiences, whereas the subjective perspective has access only to the content (better: phenomenal character) of consciousness. Hurley develops this objection by contrasting two subjects of experience, S_1 and S_2 . S_1 has two unified experiences, an experience of red at a particular location in her visual field (v_1) and an experience of hearing a violin (a_1). S_2 , however, has a partially unified consciousness. She has three experiences: an experience of red (v_1) and *two* experiences of hearing a violin (a_1 and a_2). Each of S_2 's auditory experiences are unified with her visual experience but they are not unified with each other. Importantly, a_1 and a_2 are what I call *phenomenal duplicates*: they have exactly the same phenomenal character. Hurley suggests, plausibly enough, that there would be no subjective contrast between S_1 and S_2 , and hence that there is no such thing as a distinctive partially unified phenomenal perspective. The upshot, she argues, is that objections to partial unity based on projectability cannot succeed. 'Suppositions about the structure of consciousness... are not necessarily captured by or subject to the "what it is like" test' (Hurley 1998: 166).

The first thing to note is that the objection requires S_2 to have phenomenal duplicates. We can see this by contrasting the two other subjects of experience— S_3 and S_4 . These two subjects enjoy tokens of the same experiential types: an experience of red at a particular location in the visual field, an experience of hearing a violin, and an experience of a trumpet. In S_3 's case these three experiences are fully unified with each other, but in S_4 's case the two auditory experiences are unified with the visual experience but not with each other. In this scenario, there would clearly be a subjective, introspectable difference between what it's like to be S_3 and what it's like to be S_4 . Whereas S_3

¹¹ There are deep connections between Hurley's objection and Tye's treatment of the unity of consciousness (see in particular his transparency objection to the mereological account of phenomenal unity). Both Hurley and Tye take there to be a tension between the claim that the unity of consciousness is subjectively accessible and the claim that it is a relation between token experiences. However, they respond to this tension in very different ways. Hurley retains the token experience conception of unity but rejects subjective accessibility, whereas Tye retains subjective accessibility but rejects the token experience conception of unity. The tripartite conception of experiences shows how these two claims can be reconciled with each other: the unity of consciousness is a relation between token experiences that is also subjectively accessible.

¹² I have changed the details of Hurley's case, but the changes do not alter the essential structure of her argument.

has a conjoint experience of the violin and the trumpet, S_4 has no such experience. So, Hurley's objection to the projectability argument poses a threat only if phenomenal duplicates are possible. Are they?

In §2.2 I suggested that we should individuate experiences in tripartite terms. This conception of experiences leaves no room for phenomenal duplicates, for by definition phenomenal duplicates are experiential states with the same phenomenal character that are had by the same subject at the same time. Of course, there are other accounts of how to individuate experiences, and some of those accounts might recognize the coherence of phenomenal duplication. But although there may be views of consciousness that provide Hurley's objection with a toehold, they are not views that I regard as particularly attractive. I conclude that premise (1) is secure: if partial unity were possible, then there would be such a thing as a distinctive partially unified subjective perspective.¹³

What about the second premiss—the claim that we are unable to project ourselves into a partially unified perspective? On the face of things, one might have thought that the task of projecting oneself into a partially unified phenomenal perspective poses no difficulty whatsoever. Lockwood has suggested that one could project oneself into a partially unified perspective simply by imagining each of the subject's overlapping experiences ($e_1 \& e_3$ and $e_2 \& e_3$) in turn (1989: 92). Since these two acts of imagination are successive they need not bring with them an experience of $e_1 \& e_2$ in the way that attempting to imagine $e_1 \& e_3$ and $e_2 \& e_3$ *all at once* would.

Lockwood's proposal has not enjoyed a warm reception. Peacocke worries that successive acts of imagination would fail to capture the fact that one is imagining the simultaneous experiences of a single subject (Peacocke 1994: xx), while Dainton objects that successive acts of imagination would fail to ensure that the e_3 -type experience in $e_1 \& e_3$ was numerically identical to the e_3 -type experience that occurs in $e_2 \& e_3$ (Dainton 2006: 98). At first sight these objections are puzzling. Can't one simply *stipulate* that one is imagining the simultaneous experiences of a single subject, or that the e_3 -type experience in $e_1 \& e_3$ is identical to that in the $e_2 \& e_3$ experience? Aren't the objects of our imagination up to us? As Wittgenstein pointed out, 'if a man says "I am

¹³ As an aside, note that Hurley's objection raises questions about just how to understand her notion of phenomenal unity ('co-consciousness', in her terminology). Despite claiming that the objection shows that we cannot think of the unity of consciousness in 'what it's like' terms, Hurley herself appeals to just such a view in introducing the very notion of co-consciousness. 'There is unity of consciousness when various states and their contents are co-conscious: are together within one consciousness, rather than featuring in separate conscious points of view' (1998: 88). What might it mean for conscious states to be 'together within one consciousness' if not that the states share a conjoint phenomenal character? Given Hurley's rejection of 'what it's like' accounts of the unity of consciousness it's not entirely clear what she means by 'co-consciousness'.

imagining King's College on fire", it seems absurd for his friend to respond, "How do you know it is King's College you are imagining? Are you aware that there is a precise replica of it on a film set in Los Angeles? May you not be imagining *it*?" (Wittgenstein 1965: 39). In short, it is unclear why we can't imagine what it would be like to be a partially unified subject merely by engaging in successive acts of imagination.

In order to get a better fix on this issue we need to examine the notion of projective imagination in more depth. As I am conceiving of it here, projection involves imagining something 'from the inside'. Nagel has called this state 'sympathetic imagination'. To imagine something sympathetically one 'puts oneself in a conscious state resembling the thing itself' (Nagel 1974: fn 11). It follows from this that one can sympathetically imagine only conscious states—more accurately, that one can sympathetically imagine only what it is like to *be* in a certain type of conscious state. In this respect sympathetic imagination must be strictly distinguished from perceptual and propositional imagination. Unlike acts of perceptual or propositional imagination, acts of sympathetic imagination must have the same structure as their targets, for sympathetic imagination involves *replicating* the state being imagined. One cannot sympathetically imagine an atemporal phenomenal perspective without one's own conscious state taking on an atemporal character. Nor, to use Nagel's own example, can one sympathetically imagine what it would be like to echolocate without actually having that ability (or at least something very much like it) oneself. We can now see why it is not possible to stipulate either that one is imagining the simultaneous experiences of a single subject, or that the e_3 -type experience in $e_1&e_3$ is identical to that in the $e_2&e_3$ experience, for such stipulation doesn't extend to sympathetic imagination. Engaging in successive acts of projection would only enable one to imagine being a subject with one kind of phenomenal state that was followed by another. In order to project oneself into the phenomenal perspective of a subject at a particular time one would need to engage in a *single* act of imaginative projection. In short, premiss (2) of the projectability argument appears to be secure when suitably understood.

But although both premisses (1) and (2) are plausible when suitably understood premiss (3) is not, for neither unimaginability in general nor sympathetic unimaginability in particular is a reliable guide to impossibility. The fact that we are unable to imagine what it is like to be a bat provides no reason at all for thinking that there is nothing it is like to be a bat. Nor, to take another well-worn example, does the fact that Mary in her black and white room cannot imagine what it is like to experience red give her any reason to doubt the possibility of such experiences (Jackson 1982). Premise (3) could be defended only if there were reason to think that our projective abilities limn the space of

phenomenal possibility, and there isn't. (In fact, there is good reason to think that our projective capacities extend to only a very small region of that space.) Being *able* to project oneself into the perspective of a disunified subject would be relevant to the current debate—it would be an existence proof that disunified subjects are possible, for only a disunified subject would be able to project itself into the perspective of a disunified subject—but being *unable* to project oneself into such a perspective doesn't go any way towards showing that such subjects are impossible.

The projectability argument fails, but perhaps something can be salvaged from its wreckage. Why is it so plausible to suppose that there are possible phenomenal perspectives that are imaginatively inaccessible to us? Surely our basis for thinking that such perspectives are possible is grounded in the fact that we can *conceive* of them. This, in turn, suggests that it might be useful to consider whether there might be something *incoherent* in the notion of partial unity. Perhaps we can construct a conceivability-based argument against partial unity.

Following van Cleve (1983), let us distinguish two types of inconceivability: *strong inconceivability* and *weak inconceivability*. A scenario is strongly inconceivable for *S* when *S* seems to see that it is impossible, whereas a scenario is weakly inconceivable for *S* when *S* cannot see that it is possible. As the terms suggest, strong inconceivability entails weak inconceivability but not *vice versa*: that is, if a scenario is strongly inconceivable for someone then it is also weakly inconceivable for them, but the converse conditional does not hold. (Note that 'strength' here concerns the content of the inconceivability intuition rather than its epistemic status or degree of subjective certainty.)

Examples of strong inconceivability—or at least putative strong inconceivability—are not hard to find. Scratch a putative metaphysical truth and one finds a claim of strong inconceivability underneath it. Arguably, uncaused events are inconceivable; physical objects that can survive the process of being shrunk to an extensionless point are inconceivable; square circles are inconceivable; tigers that lick all and only all tigers that do not lick themselves are inconceivable, and so on. As an example of a proposition that is weakly inconceivable consider Goldbach's conjecture that every even integer greater than 3 is the sum of two prime numbers. I take it that you neither see that Goldbach's conjecture is true nor do you see that it is false.

Is partial unity *strongly* inconceivable? That seems unlikely. Transitivity may be a deep feature of synchronic unity, but I very much doubt whether it reveals itself to us as such. Indeed, I am not convinced that first-person acquaintance with consciousness reveals *any* substantive features of it to be necessary. How could first-person reflection on experience provide one with grounds for thinking that certain features of one's own consciousness derive from the

essential nature of consciousness itself as opposed to contingent aspects of one's own cognitive architecture? There may be some sense in which partial unity is inconceivable, but it is not, I think, strongly inconceivable.

That leaves weak inconceivability. Is partial unity weakly inconceivable? I am inclined to think that it is. At any rate, I cannot 'see' that partial unity is possible, and as far as I can tell I am not alone in this. Of course, to claim that partial unity is weakly inconceivable is not to claim a great deal, for an inability to see that *P* is impossible may be due to nothing more than a cognitive limitation on the part of the cognizer. (Consider again Goldbach's conjecture.) The prudent position, it seems to me, is to retain partial unity as a potential model of consciousness, albeit one that is surrounded by a significant degree of suspicion.

Some theorists might be inclined to think that I am being overly cautious here. One might argue that reflection on the phenomenal continuity of consciousness—its unity across time—gives us positive reason to think that partial unity is coherent. As we noted above, there is good reason to think that transitivity can fail with respect to successive experiences. In hearing the notes *Do*, *Re*, and *Me*, one's experiences of *Do* might be unified with one's experience of *Re*, and one's experience of *Re* might in turn be with one's experience of *Me* without one's experience of *Do* being unified with one's experience of *Me*. But—a critic might press—if transitivity can fail when it comes to conscious states that are spread out across time, why can't it also fail when it comes to simultaneous states?

Here is not the place to enter into a protracted discussion of the mysteries of temporal experience, but I do want to briefly present two possible lines of response to the objection. First, there is at least one model of phenomenal continuity on which transitivity holds. Suppose that hearing the notes *Do*, *Re*, and *Me* involves successive punctate experiences, e_1 and e_2 , which have as their contents $\langle Do \text{ at } t_1; Re \text{ at } t_2 \rangle$ and $\langle Re \text{ at } t_2; Me \text{ at } t_3 \rangle$. Moreover, suppose that these experiences are not phenomenally unified with each other: although both e_1 and e_2 represent $\langle Re \text{ at } t_2 \rangle$, the representation of *Re* that occurs in the context of e_1 is completely distinct from that which occurs in the context of e_2 . (As an aid to intuition, we might suppose that visual experience is pulse-like, and there is a very short temporal gap between e_1 and e_2 .) This view appears to account for temporal experience without supposing that there is any failure of transitivity for phenomenal unity across time. This account of phenomenal continuity may not be particularly plausible, but it does undermine the thought that we *must* give up on transitivity in order to account for phenomenal continuity.

A second—and perhaps more compelling—response to the objection is that there may be constraints on the unity of consciousness that don't apply to its continuity. In fact, one might venture the thought that the unity of consciousness

is in part distinguished from its continuity by virtue of the fact that transitivity must hold with respect to simultaneous experiences but not with respect to successive experiences. We can see how experiences of *Do* and *Me* might each be unified with an experience of *Re* without being unified with each other, for the temporal locations of these experiences generates a gap between the experiences of *Do* and *Me* that is not matched by a corresponding gap between the experiences of *Do* and *Re* or between the experiences of *Re* and *Me*. There are, of course, no temporal gaps between simultaneous experiences, and so any failure of transitivity here would need to derive from another source. There may be such sources—for example, there might be something in the structure of the neuro-functional relations underlying consciousness which allow for partial unity—but the central point is that there is no inconsistency in allowing that transitivity can fail for successive experiences but at the same time asserting that it must hold for simultaneous experiences.

2.5 Conclusion

Although the unity of consciousness could hardly be a more familiar phenomenon we are still some way from having an adequate theoretical grasp of what it is for consciousness to be unified. The focus of this chapter has been to develop and defend a mereological account of phenomenal unity, according to which experiences are phenomenally unified with each other exactly when they occur as parts of a single experience. Section §2.2 focused on the notion of experience that lies at the bottom of the mereological account. I began by noting that the mereological account is inconsistent with Tye's one-experience view, according to which experiences are to be identified with entire streams of consciousness. However, I suggested that there is little reason to embrace this revisionary conception of experience, and that we can quite happily think of certain components of a stream of consciousness—what Tye calls 'experience stages'—as experiences in their own right. But not every stretch of a stream of consciousness is an experience. I suggested that we should think of experiences in terms of the instantiation of phenomenal properties, where a phenomenal property is a property that there is 'something it is like' to enjoy. More precisely, I advanced a tripartite conception of experiences: a token experience can be equated with the instantiation of a phenomenal property by a subject at a time.

In §2.3 I addressed a number of objections to the mereological account of phenomenal unity. Some of these objections were well targeted but could be

answered from within the mereological framework (the self-undermining objection; the phenomenal bloat objection; the transparency objection); others missed their target, either because they assumed that the mereological account has commitments that it doesn't (the building block objection), or because they assumed that it is trying to capture a kind of unity that it isn't (the Jamesian objection).

In the final section of this chapter I turned from the mereological analysis of phenomenal unity as such to a question that it raises: is it possible for a set of simultaneous experiences to be merely partially unified? The central objection to the coherence of partial unity took the form of the projectability argument. Although influential criticisms of the projectability argument turned out to be unconvincing, I argued that the objection nonetheless fails because it presupposes an overly tight connection between projectability and possibility. We saw however that there is something to be salvaged from the demise of the projectability argument: an inconceivability argument. We may not be able to see that partial unity is impossible, but we cannot—I suggest—see that it is possible. The upshot of these reflections, I suggest, is that an air of suspicion ought to accompany any talk of partial unity. As with many of the other issues examined in this chapter, we will return to the question of partial unity in later chapters.