10

The Quilt of Consciousness

Whether spun from the fabric of philosophical imagination or patched together on the basis of fMRI images, few theories of consciousness are constructed with the unity of consciousness in mind. This is rather odd, for if consciousness is unified—as I have argued it is—then our models of consciousness ought surely to be informed by this fact. Any adequate account of consciousness must be consistent with, and perhaps even explain, the unity of consciousness. In this chapter I examine theories of consciousness with respect to how well they might accommodate this constraint. In order to dampen down expectations, let me say from the outset that I am not in the business of solving 'the hard problem' (Chalmers 1996) or closing 'the explanatory gap' (Levine 1983). I am inclined to think that explaining how to get experience from neural activity is not unlike trying to get 'numbers from biscuits or ethics from rhubarb' (McGinn 1993: 155), and neither of those activities strike me as useful investments of either my time or yours. Instead, my aim is the rather more modest one of sketching a framework that might inform theory-building in consciousness studies. My hope is to dislodge some bad ideas, and put some better ones in their place.

10.1 Atomism versus holism

One way to distinguish between theories of consciousness is in terms of how they conceive of the structure of the phenomenal field. Let us distinguish two general orientations that can be adopted here: an *atomistic* orientation and a *holistic* orientation. 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.[1]

The contrast between atomism and holism is no mere formal nicety but has important methodological implications. Atomists will be inclined to recommend a 'bottom-up' approach to the study of consciousness according to which we should focus on understanding the mechanisms responsible for generating the atoms of consciousness. Holists, by contrast, will have deep reservations about this approach. They will doubt whether we can understand consciousness by focusing on the components of the phenomenal field in a piecemeal manner. Instead, the holist will recommend a top-down methodology, according to which we should look for the mechanisms implicated in the construction of the entire phenomenal field.

The distinction between atomism and holism cuts across many of the standard ways of classifying theories of consciousness. Let us consider how the distinction might play out within the context of three broad approaches to consciousness: monitoring approaches, functionalist approaches, and neurally based approaches.

Monitoring accounts of consciousness hold that a mental state is conscious in virtue of being monitored in some way or another.[2] Discussion of this approach has focused on the nature of the monitoring representation ('Is it thought-like or perception-like?') and on the relationship between the monitored state and the monitoring state ('Might the monitoring state be identical to the monitored state or must they be distinct?'). In contrast, there has been rather little discussion of whether monitoring accounts of consciousness should be developed in atomistic terms or holistic terms. The monitoring approach as such can be developed in either direction, and in fact both atomistic and holistic versions of the monitoring approach can be found in the literature. Rosenthal (2003) defends an atomistic version of the monitoring approach, although his account has holistic elements insofar as he holds that the monitoring thoughts responsible for consciousness 'clump together' sets of monitored states. Van Gulick, on the other hand, defends a holistic version of the monitoring approach, according to which the 'transformation from unconscious to conscious state is not merely a matter of directing a separate and distinct meta-state onto the lower-order state but of "recruiting" it into the globally integrated state that is the

[1] This distinction has something in common with Shoemaker's (2003) distinction between atomistic and holistic theories of consciousness and even more in common with Searle's (2000) distinction between building block and unified field models of consciousness, but neither author quite captures the contrast that I am after here.

[2] For important collections of papers on monitoring accounts of consciousness see Gennaro (2004) and Kriegel & Williford (2006).

momentary realization of the agent's shifting transient conscious awareness' (van Gulick 2004: 74f.).

The functionalist approach to consciousness can also be developed in either atomistic or holistic directions. At the heart of functionalist accounts is the claim that a state is conscious in virtue of the functional role that it plays within the subject's cognitive economy, where this role is typically identified with some kind of capacity for broad cognitive and behavioural control. Much of the discussion of functionalism has focused on its epistemic status ('Is it an *a priori* truth or an *a posteriori* one?') and on the question of how to best characterize the functional role distinctive of conscious states ('What exactly does broad cognitive and behavioural control amount to?'), but there has been rather less discussion of the implications of functionalism for accounts of the structure of consciousness.

Many of the most influential versions of functionalism are most naturally understood in atomistic terms. Consider, for example, Dretske's (1995) account of consciousness, according to which phenomenally conscious mental states are identified with non-conceptual representations that supply information to a cognitive system for calibration and use in the control and regulation of behaviour (Dretske 1995: 19). On this account there is good reason to suppose that (say) a visual state and affective state might qualify as independent atoms of consciousness, for the two states are likely to control and regulate behaviour independently of each other. Consider also Tye's (1995) PANIC model of consciousness. This account qualifies as a version of functionalism in virtue of the fact that the 'P' in his acronym stands for 'poised', where a state is poised exactly when its content stands ready and in position to make a direct impact on the subject's belief/desire system (1995: 138). Tye's account of consciousness is *formally* holistic for he holds that the only states that qualify as experiences are entire streams of consciousness (see §2.1), but it is atomistic in spirit, for there is nothing in the account which ensures that various non-conceptual representations will be poised to control thought and action conjointly—that is, as the components of a single state. Other versions of functionalism are more naturally understood in holistic terms. For example, Shoemaker claims that 'the factors that go into making a particular mental state conscious are inextricably intertwined with those that go into making different states 'co-conscious', i.e. go into constituting a unified state of consciousness of which that state is a part' (2003: 58). Although Shoemaker's account falls short of a full-blown commitment to holism, it clearly has more of a holistic orientation than do the versions of functionalism developed by Dretske and Tye.

Finally, let us examine how the contrast between atomism and holism plays out in the context of neurally-inspired accounts of consciousness. A prominent

(and indeed radical) version of atomism is defended by Zeki and Bartels (1999). They argue that activity at each node of a processing–perceptual system generates its own 'micro-consciousness'. These micro-consciousnesses are independent units of consciousness that involve only the registration of fine-grained conscious features—colour, motion, shape, and so on. A less extreme version of atomism appears to be implicit in Lamme's recurrent processing account of consciousness, according to which visual experience involves recurrent (or re-entrant) processing between lower and higher regions within the visual cortex (Lamme 2006; Lamme & Roelfsema 2000). Although Lamme's focus is on visual experience, the recurrent processing approach can be applied to consciousness more generally. When so extended it appears to generate an atomistic conception of consciousness, for there is nothing in the approach itself which limits the subject to one stream of recurrent processing at a time, or which requires multiple streams—streams located within different perceptual systems—to be bound together into a single recurrent cascade.

Other neurally based models of consciousness are more naturally understood in holistic terms. Consider the global workspace approach to consciousness, originally proposed by Baars (1988) and developed more recently by Dehaene and Naccache (2001). Global workspace accounts hold that 'conscious experience emerges from a nervous system in which multiple input processors compete for access to a broadcasting capability; the winning processor can disseminate its information globally through the brain' (Baars 1993: 282). This approach suggests a holistic picture of consciousness, insofar as its advocates typically insist that entry into the workspace is gated by a 'winner takes all' competition, with only a single representational state admitted to the workspace at any one point in time (see e.g. Dehaene & Changeux 2004).

We have seen that the contrast between atomistic and holistic approaches to consciousness cuts across many of the established classifications for theories of consciousness. My interest in this chapter is not with the question of whether theories of consciousness should be pursued within (say) a monitoring framework or a neural framework, but whether we should think of consciousness in holistic terms or atomistic terms. It will come as little surprise that I will argue in favour of holism. My case against atomism begins in §10.2 and §10.3, where I argue that various aspects of the representational unity of consciousness put pressure on some of the more extreme manifestations of atomism. However, these sections are best thought of as entrées to the main business of this chapter, which takes place in §10.4. There, I argue that atomistic accounts of consciousness are at odds with the unity thesis, and conclude that we should adopt a holistic conception of the structure of consciousness. §10.5 identifies and attempts to disarm a number of objections to holism, and §10.6 presents a

sketch of what a plausible version of holism might look like. But before we turn to the unity thesis and its implications, we need to first examine possible constraints imposed by the representational unity of consciousness.

10.2 Binding

In a special issue of *Neuron* dedicated to the binding problem Adina Roskies asks, 'Will the solution to the binding problem be the solution to the mystery of consciousness?' (Roskies 1999: 9). Roskies has not been alone in suggesting that the connections between consciousness and the binding problem run deep, and it is widely held that a model of binding might shed essential light on the nature of consciousness itself.[3] Although I doubt that binding will unlock the puzzle of consciousness, I do think that it provides us with reason to reject certain radical versions of atomism.

Let us start with binding. Most generally, binding is the process of bringing information together—synthesizing it, as Kant would say. As Kant also noted, binding takes a number of forms. One form of binding occurs when one applies a concept on the basis of perceptual experience, and in so doing integrates conceptual and perceptual representations. Binding of another sort takes place when one integrates a sequence of perceptual experiences, and builds up a global representation of one's environment. But the kind of binding that has been the focus of 'the binding problem' is *feature binding*: the integration of perceptual features—colour, shape, texture, identity, and so on—into coherent percepts of objects. Feature binding is essential to object unity (Chapter 1).

Because distinct types of perceptual features are processed in different locations within cognitive architecture, there is a question about how the brain binds these features together so as to create representations of integrated perceptual objects. In fact, there are really two feature binding problems: a *selection* problem and a *tagging* problem (Robertson 2003; Treisman 1996, 2003). Suppose that you are looking at a grey donkey carrying a child who is wearing a blue jacket. Your visual system needs to do two things here. First, it needs to select the right features to put together: it needs to represent the donkey as grey and the jacket as blue rather than vice versa. Secondly, it needs to make sure that those features which have been selected are 'tagged'—that is, bound together as a functional unit and made available as such to downstream systems. Of most interest to us here is the tagging component of the binding problem.

[3] See e.g. Engel et al. (1999a); Revonsuo (1999); Sauvé (1999).

Some theorists have suggested that the binding problem is a pseudo-problem, a relic of a long-discredited commitment to some kind of Cartesianism.

> Those for whom the consciousness module is a lost cause appeal to so-called 'binding', although they do see it as a 'problem'. Because the attributes of a perceived object appear bound, there must be some brain-glue (synchrony, co-oscillation?), that conjoins the separate and scattered representations of its attributes into the object representation . . . This would make sense if there were an inner observer for whose viewing pleasure the screens of the Cartesian Theatre display all this binding. But the brain is unsupervised and uncentered, and no one is watching. (Kinsbourne 2000: 546)

Kinsbourne is right to point out that there is no homunculus inspecting images projected onto the screen of subjectivity, but the binding problem survives the closure of the Cartesian Theatre. No matter how unsupervised and uncentred neural functioning might be, we still need an account of how distributed features are bound together to form representations of unified objects, and how—once bound—they are kept together so as to exert cognitive and behavioural influence as a unit.

So much for binding itself—why think that binding has anything to do with consciousness? As best I can tell, there are two reasons why theorists have been attracted to the thought binding might bear on the analysis of consciousness. The first of these reasons concerns the influence of temporal synchrony accounts of binding. Synchrony accounts of binding hold that the tagging of perceptual features is achieved by the phase-locked activation of distributed neural areas.[4] The temporal synchrony model replaces so-called 'grandmother cells' (that is, cells that fire only in response to particular objects such as one's grandmother) with clusters of neurons that respond to specific, 'low-level' features.[5] Suppose that one is viewing a scene containing a red square and a green circle. The idea is that the neurons that code for redness and squareness would fire in synchrony with each other as would the neurons that code for green and circularity, but the first sets of neurons would fire out of phase with the second set of neurons. However, temporal synchrony would provide a bridge between binding and consciousness only if there were evidence that synchronization is involved in both feature binding and consciousness. There is indeed some evidence in

[4] For representative defences of the temporal synchrony approach to binding see Milner (1974); von der Malsburg (1995); Eckhorn (1999); Engel et al. (1999a); Engel & Singer (2001); Roeflsema et al. (1997); Singer (2009); Singer & Gray (1995). For critical discussion of the approach see Gold (1999); O'Reilly et al. (2003); Shadlen & Movshon (1999).

[5] This account of binding is sometimes referred to as the '40Hz hypothesis' because the relevant synchrony is thought to occur at around 40Hz (in fact, between 35 and 70Hz). However this label is somewhat misleading in that it suggests that the hypothesis requires synchronized oscillations, but in fact non-oscillatory signals can also be synchronized (Engel et al. (1999b)).

favour of synchrony models of binding (although its force is highly disputed), but I know of no evidence implicating synchrony in the explanation of consciousness. In fact, the only reason to think that temporal synchrony might be implicated in models of consciousness involves appealing to a prior link between binding and consciousness, which is precisely what is in question at this point.

The second motivation for thinking that the mechanisms of binding and those of consciousness might be deeply connected derives from the fact binding and consciousness are robustly *correlated*. This thought has two components: feature binding occurs only in the context of conscious representations, and consciousness doesn't harbour unbound features. This correlation wouldn't demonstrate that there is an explanatory or constitutive connection between binding and consciousness, but it would certainly provide that view with some support. But is there a robust correlation between consciousness and binding? Let us begin with the question of whether binding requires consciousness. Is it the case that features must be conscious in order to be bound together, or can binding occur outside of consciousness?

Although influential discussions have assumed that consciousness is required for feature binding (Crick & Koch 1990a, 1990b), it is not difficult to find examples of feature binding outside of consciousness. Patients with blindsight appear to be capable of binding together primitive features in their blind field (Kentridge et al. 1999, 2004). Certain forms of priming also suggest that feature binding can occur independently of consciousness. In order for the word 'dog' to function as a semantic prime one must bind together its parts so as to form a representation of the word 'dog' as such (Neumann & Klotz 1994; Marcel 1983). Consider also the kind of priming that occurs in unilateral neglect (Ladavas et al. 1993; Farah 1997). Processing of a word presented in the right visual field can be facilitated by the brief presentation of an associated word in the neglected left visual field, even though the neglected word appears not to enter consciousness. In order for the neglected word to facilitate right visual field processing it needs to be represented as such, and this of course requires the binding of its constituent features. A further form of visual binding that occurs outside of consciousness concerns the integration of (unconscious) dorsal stream with (conscious) ventral stream representations. In short, there is ample reason to think that feature binding is not restricted to consciousness.

Consciousness may not be necessary for feature binding, but perhaps feature binding is necessary for consciousness. In fact, it is probably this claim that represents the most influential conception of the link between binding and consciousness:

> Conscious access in perception is always to bound objects and events . . . Experienced objects have colours, locations, orientations. They may not always be correctly bound;

in fact, when we first look at a complex multiobject scene they probably are not. But it seems impossible to even imagine free-floating shapes, colours, or sizes. Using language or other symbols, we can abstract particular properties (a kind of unbinding), but this is not part of our perceptual experience. (Treisman 2003: 97; see also Engel et al. 1999a, 1999b)

According to what we might call the *binding constraint*, perceptual features cannot enter consciousness without first being bound together in the form of objects. If true the binding constraint would demonstrate that the mechanisms of consciousness must be 'downstream' of the mechanisms responsible for feature binding. It would also have implications for the size of the minimal units of perceptual consciousness.

Is the binding constraint true? It is certainly extremely plausible. However, it is difficult to be more definitive than that without an account of what qualifies as a perceptual feature. We may have a rough and ready idea of what the features of visual experience are—although even here there is ample room for debate—but our grip on what counts as a feature is much less secure once we leave the domain of vision. What are the feature of auditory experience? Are they pitch, volume, and timbre? What are the 'features' of olfactory experience? And what about other forms of sensory consciousness, such as pain? Is the aversive character of a pain one feature, its felt bodily location another, and its qualitative character—its 'painfulness'—a third? These questions cannot be answered by appeals to introspection but require detailed accounts of the architecture of perception. Pending the development of such accounts any definitive verdict on the tenability of the binding constraint would be premature.

What we can say, however, is that there is at least a *prima facie* case for thinking that certain kinds of sensory features can enter consciousness without first being bound. Consider experiences of 'free-floating' anxiety. As best I can tell such features are not bound to anything. They are not bound to particular body parts, nor are they bound to one's body as a whole. Consider also olfaction. Whatever features olfactory experience might involve, it is far from clear that they must be bound together in the form of object- or event-involving percepts in order to enter consciousness. The fact that olfactory experience lacks the rich spatial structure of (say) visual experience raises questions about just what a binding-constraint for olfaction might even look like.

But even if certain kinds of sensory features can enter consciousness in an unbound form, feature binding can still provide us with a robust constraint on theories of consciousness, for there is good reason to think that perceptual features of most kinds are conscious only in the context of bound percepts.

And if this is right, then certain radical forms of atomism—such as Zeki and Bartels's 'micro-consciousness' approach—must be false. According to Zeki and Bartels, the units of visual consciousness can be identified with the perceptual features—colour, motion, shape, and so on—that are processed at the various 'nodes' of the visual system. But if that were so, then these features ought to be capable of independent existence within the stream of consciousness, and that doesn't seem to be the case. Contrary to what Zeki and Bartels suggest, we should not think of feature binding as a process of 'bringing different conscious experiences together' (Bartels & Zeki 1998: 2330), for there is in general no perceptual experience prior to feature binding. Binding might not be the sword that severs the Gordian knot of consciousness, but it is surely a form of conscious unity that no theorist can afford to ignore.

10.3 Inter-sensory integration

The binding constraint places pressure on some of the more extreme manifestations of atomism but it is not a particularly challenging constraint to meet and few versions of atomism fall foul of it. More challenging is a constraint to be explored in this section—a constraint of representational integration.

The commonplace distinction between the senses encourages us to think of the stream of experiences as divided into modality-specific chunks. We naturally fall into the temptation of referring to visual experiences, auditory experiences, and so on. Now, although there is a perfectly legitimate sense in which we can think of the stream of consciousness as containing modality-specific segments it would be a mistake to think of such states as atoms of consciousness.

We can see why by considering the phenomenon of inter-modal integration.[6] There is now a huge literature devoted to the complex and extensive ways in which the contents of experiences within one modality are intimately dependent on those in another modality. Consider first the temporal content of experience. In what is known as 'temporal ventriloquism', the subsequent presentation of an auditory stimulus changes the perceived temporal location of a previously presented visual stimulus (Morein-Zamir et al. 2003; see also Kamitani & Shimojo 2001). In fact, following a flash with a sound moves the perceived time of the flash back whilst following it with another sound moves it

[6] For reviews of the empirical literature see Calvert et al. (2004); Driver & Noesselt (2008); Ernst & Bülthoff (2004); Lalanne & Lorenceau (2004); Macaluso & Driver (2005); Spence & Driver (2004); Spence and Squire (2003); Spence et al. forthcoming; Stein & Meredith (1993); Stein et al. (2009). For reflections on the theoretical implications of this research see Shimojo & Shams (2001); O'Callaghan (2008) and forthcoming.

forward in time (Fendrich & Corballis 2001). In another example of the intermodal dependence of temporal content, the perceived rate of visual flicker can be modulated by the rate at which a concurrent stream of auditory input is presented (Recanzone 2003; McDonald et al. 2005). The window within which the temporal structure of visual and auditory experience is integrated is not fixed but dynamic, its parameters change in order to accommodate the fact that sound increasingly lags behind vision as the source of the sound increases in distance (Sugita & Suzuki 2003).

Generally speaking, the temporal content of audition modulates that of vision rather than vice versa, but the converse is the case when it comes to the spatial content of perceptual experience. Here visual information generally dominates that of other perceptual modalities, such as audition (Howard & Templeton 1966; Bertelson 1999) and touch (Pavani et al. 2000; Rock & Victor 1964). Indeed, the influence of the spatial content of vision on other modalities can be quite profound. For example, the direction of visual motion can actually *reverse* the experienced direction of an auditory stream (Zapparoli & Reatto 1969; Soto-Faraco et al. 2004).

Intermodal effects are not limited to the spatial and temporal contents of experience but can even modulate the *number* of objects that subjects experience. In what is known as the auditory-flash illusion, subjects will misperceive a single flash of light as two flashes when it is paired with two beeps (Shams et al. 2000). Intermodal effects are not limited to properties that are available via more than one modality (the common sensibles) but can also be found for properties that are proprietary to a particular sense, such as colour. For example, tactile stimulation at a particular location can improve the perceptual discrimination of colour at that and nearby locations.[7]

Intermodal integration also occurs for high-level, categorical information. Perhaps the most well known of such effects is the McGurk effect, in which dubbing the phoneme /ba/ onto the lip movements for /ga/, produces, in normal adults, an auditory percept of the phoneme /da/ (McGurk & McDonald 1976). More recently, Chen and Spence (2010) have shown that semantically congruent auditory input can help subjects identify masked pictures. Subjects who had been played the sound of a barking dog were more likely to recognize a masked picture as that of a dog than were controls who had been exposed only to white noise.

Is the extent of inter-perceptual integration surprising? Perhaps. We might have assumed that a more reliable architecture would have entrusted the job of

[7] Spence et al. (2004); see also Frassinetti et al. (2002); Lovelace et al. (2003); McDonald et al. (2000).

adjudicating inter-sensory conflict to post-conscious executive systems. But this is not the cognitive architecture that we have inherited. Instead, the brain resolves most inter-sensory disputes prior to consciousness on the basis of '*a priori*' assumptions about the relative reliabilities of the various senses in particular contexts, with the result that 'the conscious subject' is unaware of the very existence of the disputes.[8] But although this arrangement might seem surprising, a moment's reflection reveals its advantages: eliminating inconsistencies prior to consciousness frees post-conscious mechanisms up for other tasks. (And of course some inconsistency does make it through to consciousness, as we noted in Chapter 3.)

The ubiquity of inter-modal integration puts further pressure on atomistic approaches to consciousness. Not only does it provide additional evidence against the view that perceptual features qualify as atoms of consciousness, it also tells against more moderate forms of atomism that conceive of the units of consciousness in modality-specific terms. Consider, for example, O'Brien and Opie's endorsement of the atomistic view:

> our instantaneous phenomenal experience is a complex amalgam of distinct and separable conscious events; not a serial stream, but a mass of tributaries running in parallel . . . a conscious individual does not have a 'single consciousness,' but several distinct phenomenal consciousnesses, at least one for each of the senses, running in parallel. (O'Brien & Opie 1998: 387 see also O'Brien & Opie 2000.)

It is not obvious how this conception of consciousness might be squared with inter-modal integration. How could we account for the subtle interplay between the contents of the various modalities if the senses are a 'mass of tributaries running in parallel'? Inter-modal integration suggests that the stream of perceptual experience is best thought of as highly braided rather than as composed of sense-specific tributaries that generate experience in splendid isolation from each other. The senses are not hermetically sealed off from each other, but function as highly interdependent channels. A subject's perceptual experience in any one 'modality' is the result of complex interactions between any number of sensory channels, and the hope that one might be able to identify stable, modality-specific mechanisms underlying it is, I suggest, a vain one. Inter-sensory integration does not itself show that the atomistic conception of consciousness is untenable, but it does suggest that whatever atoms of consciousness there might be are unlikely to take the form of modality-specific chunks.

[8] See Alais & Burr (2004); Ernst & Bülthoff (2004); and Helbig & Ernst (2007).

10.4 Implications of the unity thesis

Although both the binding constraint and intermodal integration put pressure on certain manifestations of atomism, it would be stretching things to say that they are at odds with atomism as such. But let me turn now to a unity constraint that does put pressure on all forms of atomism: the unity thesis.

Atomists hold that the subject's total phenomenal state is built up out of units of consciousness. As we have noted, there are a number of ways in which atomists might conceive of the 'size' of these units. Radical atomists will hold that the atoms of consciousness are perceptual features. More moderate forms of atomism will suggest that the atoms of consciousness can be thought of in terms of modality-specific experiences, such as integrated visual representations of scenes. Still other forms of atomism might conceive of the units of consciousness as having a structure that cuts across the traditional distinction between the senses—for example, the atomist might suggest that the units of consciousness correspond to multi-modal representations of objects. But whatever their take on the units of consciousness, atomists of all stripes must account for the fact that those units don't occur as independent elements of consciousness but as the components of an overall phenomenal field. How might atomists respond to this challenge?

One response is to simply deny that the atoms of consciousness *are* unified within the context of an overall phenomenal field. According to this kind of atomist, the so-called unity of consciousness is nothing more than illusion. Readers who have made it thus far will not be surprised that I am inclined to simply set this response to one side. In my view the only tenable response to this challenge is to posit a mechanism that might account for the fact that the atoms of consciousness are generally—if not invariably—unified with each other. This mechanism would not be responsible for consciousness as such. Instead, its role would be to ensure that the subject's conscious states are phenomenally unified with each other—subsumed by a total phenomenal state. Holists, of course, have no need to posit a mechanism that is specifically responsible for phenomenal binding, for the holist holds that the components of a subject's total phenomenal state are brought into being *as* the constituents of that state.

Although some atomists have recognized the need for phenomenal binding (see e.g. Zeki and Bartels 1999), most atomists have been curiously reluctant to discuss, let alone posit, such a mechanism. There is good reason for this reluctance, for there is little evidence of its existence. If there were such a mechanism then we would expect it to occasionally malfunction, with the result that the subject would be left with phenomenal fragments—

units of consciousness that would no longer be integrated into phenomenal wholes. But to the best of my knowledge neuropsychology furnishes us with no examples of phenomenal fragmentation. As we have seen in the preceding chapters, although there are plenty of syndromes in which other forms of the unity of consciousness break down, there are no syndromes in which the *phenomenal* unity of consciousness breaks down—at least, so I claim.

Of course, my defence of the unity thesis is open to any number of challenges. A critic might argue that my treatment of (say) the split-brain syndrome is unsatisfactory, and that there are periods during which split-brain patients have two streams of consciousness. But even if this were right it would not provide much comfort for the atomist given that the kinds of breakdowns in phenomenal unity that could conceivably characterize the split-brain hardly reveal the *pre-existing* structure of consciousness. Although there is some elasticity within the atomistic approach as to how 'big' the atoms of consciousness are, on no plausible version of the view is the typical subject's total phenomenal state composed of only two atoms, one grounded in each hemisphere. Parallel points apply to each of the other syndromes that might provide counter-examples to the unity thesis. In each case, the kinds of splits that might be thought to occur would be imposed on the structure of consciousness 'from without' rather than along pre-existing fault-lines. (The appropriate analogy for such breakdowns in phenomenal unity would be that of splitting a coconut with an axe rather than segmenting an orange into pieces). Feature binding and inter-sensory integration 'problematize' certain radical forms of atomism, but the unity thesis suggests that the entire approach is wrong-headed.

This 'unity thesis argument' is the star witness in the case against atomism, but it does not carry the case on its own. Two further features of consciousness argue in favour of holism. We might call the first of these features the *dynamic* structure of consciousness. As Koch (2004) has argued, entry into the stream of consciousness can be thought of as taking the form of a competition between coalitions of contents, each one of which struggles to make its voice heard above that of its fellows. Coalitions that win this competition—whether by top-down control or stimulus-driven attention—enter the stream of consciousness; the losers hover in the wings, waiting for their moment. The forces behind these dynamic changes are global and domain-general in nature. Whether or not a particular coalition makes its mark on consciousness depends not on its intrinsic properties but on its strength relative to those with which it is in competition. These dynamic features are hard to square with atomism. Why should entry into consciousness be dynamically gated in this way if different types of conscious states are produced by autonomous mechanisms?

Holism also receives some support from the fact that background states (or 'levels') of consciousness are domain-general. Background states of consciousness typically characterize the full spectrum of a subject's conscious states rather than some particular subset of those states. It is *the subject* that is awake, dreaming, hypnotized, delirious, and so on. We might think of background states of consciousness as ways in which the subject's overall phenomenal field is modulated. The global nature of background states of consciousness is also evident in transitions between consciousness and unconsciousness. Typically, the re-acquisition of consciousness takes place all at once rather than in (say) modality-specific stages. Consciousness may dawn gradually, but it dawns gradually over the whole. Similarly, when consciousness is lost it is typically lost 'as a whole' rather than in (say) modality-sized chunks. Although their force may be difficult to quantify, these indirect considerations provide additional reasons to embrace a holistic approach to the structure of consciousness in favour of atomism.

10.5 Correlates, causes, and counterfactuals

Although few theorists might describe themselves as card-carrying atomists, the view has wide currency within the contemporary literature. In fact, one might even describe it as a kind of orthodoxy. In this section I consider three of the most influential—and, indeed, potent—motivations for the view.

The central argument for atomism takes as its point of departure the claim that the neural mechanisms underpinning consciousness—the neural 'correlates of consciousness', as they are often described—are not to be found at any one location but are scattered throughout the brain.

> The multiplicity of cortical loci where correlations with awareness have been found provides some evidence against one of the oldest ideas about consciousness, that the contents of awareness are represented in a single unitary system. . . . Instead, the data described above seem more consistent with a view in which the contents of current awareness can be represented in many different neural structures. However, one could still argue that the neural correlates described above are not in fact the actual representations that constitute the conscious percept, but merely information that is likely to make it onto the (as-yet-undiscovered) screen of awareness, so the possibility of such a unitary awareness system is not definitively ruled out by these data. In contrast to the idea of a unitary and content-general Cartesian theatre of awareness, the data summarized above fit more naturally with the following simple hypothesis: *the neural correlates of awareness of a particular visual attribute are found in the very neural structure that perceptually analyzes that attribute*. (Kanwisher 2001: 97, emphasis in original)

Generalizing Kanwisher's comments beyond vision, we might put the challenge as follows: doesn't the fact that the neural correlates of consciousness are distributed across a multiplicity of cortical loci demonstrate that consciousness has an atomistic structure?[9]

I think not. Let us begin by noting that recent developments in cognitive neuroscience paint a rather more dynamic picture of the relationship between neural structures and the analysis of particular attributes than that suggested by Kanwisher's comments. According to a recent review, 'even classic sensory-specific areas (perhaps even primary cortices) can be influenced by multisensory interplay' (Driver & Noesselt 2008: 14). Receptive fields are rarely fixed and stable, but respond to changes in the creature's behavioural orientation. For example, eye position can modulate activity in primary auditory cortex (Fu et al. 2004; Werner-Reiss et al. 2003). In fact, many of the examples of inter-sensory integration reviewed in §10.3 are likely to result from the inter-sensory modulation of early perceptual areas.[10]

Even so, the atomist might respond, there is surely *some* sense in which the neural correlates of consciousness are localized to particular cortical areas. How might the holist account for this fact?

To fix ideas, let us consider one particular example of neural localization: the relationship between visual experiences of motion and activity in MT.[11] This relationship is clearly an intimate one. In order to selectively modulate such experiences one ought to target MT activity rather than activity in some other part of the brain. Of course, one needn't target MT if one wants to manipulate the subject's visual experience of motion but is not too particular about what other kinds of effects one might have on the subject's experience. For example, one could intervene on brain-stem systems. Not only would an intervention of this kind lead to the elimination of visual experiences of motion, it may also lead to the elimination of all forms of consciousness. These reflections show that we need an account of the neural correlates of consciousness that does justice to the

[9] Although I follow tradition and refer to the neural states underpinning consciousness as its neural *correlates*, I do not assume that these states are merely correlated with consciousness. In fact, I think it likely that they stand in a rather more intimate relation—such as realization, constitution, or even identity—to conscious states. However, I use the relatively non-committal 'correlate' in order to avoid taking a stance on the question of just how neural states and conscious states are related. See Chalmers (2000) and Hohwy (2007) for discussion of the notion of a neural correlate of consciousness.

[10] See Ghazanfar et al. (2005); Kayser & Logothetis (2007); Lakatos et al. (2007); Molholm et al. (2002); Schroeder & Foxe (2005); Senkowski et al. (2005); Watkins et al. (2006).

[11] For evidence connecting activity in MT with the visual experience of motion see Britten et al. (1992); Cowey & Walsh (2000); Heeger et al. (1999); Huk et al. (2001); Kammer (1999); Kourtzi & Kanwisher (2000); Rees et al. (2002); Théoret et al. (2002); Zihl et al. (1983).

fact that it is possible to intervene on consciousness both selectively and non-selectively. We need a conception of the neural basis of consciousness that captures the fact that MT activity is in some way essential to the visual experience of motion, but which also does justice to the fact that MT activity will generate experiences of motion only in the context of a conscious creature.

In light of the foregoing, let us distinguish between three types of neural correlates. A state's *total* neural correlate is that neural state that is *minimally sufficient* for its existence. (Note that a total neural correlate should not be confused with the neural correlate of a total phenomenal state.) We can divide a state's total correlate into two components: a *differentiating* correlate and a *non-differentiating* correlate.[12] A state's differentiating correlate is that component of its total correlate that accounts for its content. A state's non-differentiating correlate is that part of its total correlate that remains once its differentiating correlate is 'removed'. Whereas differentiating correlates *distinguish* one kind of conscious state from another, non-differentiating correlates are *shared* by a subject's conscious states. For example, visual experiences of motion and tactile experiences of one's feet will have unique differentiating correlates but common non-differentiating correlates (at least if we are considering experiences that are had by the same subject at the same time). Note that although I have contrasted differentiating correlates from non-differentiating correlates the distinction is actually a graded one, for one correlate can be more or less differentiating than another. Some non-differentiating activity might be implicated in (say) all and only visual experience; other non-differentiating activity may be implicated in all and only affective experience. I have ignored this complication in what follows in the interests of keeping the discussion manageable.

The distinction between differentiating and non-differentiating correlates is widely appreciated but not under these labels. Instead, theorists typically distinguish between 'core correlates' and 'enabling correlates' (see e.g. Chalmers 2000; Block 2005; Koch 2004). I think these terms are somewhat unhelpful, for they encourage an atomistic conception of the structure of consciousness. Referring to differentiating correlates as 'core' encourages us to view them as the fundamental mechanisms of consciousness—the systems that *really* generate consciousness—and we have seen that that assumption is problematic. Indeed, if either of these two types of correlates deserves to be thought of as 'core' it is *non-differentiating* correlates, for it is these correlates that are common to—and hence 'lie at the core of'—the subject's various conscious states. By the same token, referring to non-differentiating correlates as 'enabling correlates' down-

[12] I owe these labels to David Chalmers.

grades their status to that of the 'electrical supply'—systems whose primary job is to merely ensure that the differentiating correlates are activated. In other words, these terms encourage an atomistic picture of consciousness—a picture on which the fundamental mechanisms of consciousness are distributed across the cortex like so many lights on a Christmas tree.

Return to Kanwisher's claim that the neural correlates of awareness of a particular visual attribute are found 'in the very neural structure that perceptually analyzes that attribute.' We are now in a position to see why this claim is, at best, highly misleading. Any particular visual attribute might involve highly localized differentiating activity, but a state's differentiating correlate should not be confused with its total correlate. Consider once again the relationship between activity in MT and visual experiences of motion. Is it plausible to suppose that MT activity is not only the differentiating correlate of such experiences but is also their total correlate? Hardly. At the very least, the only evidence we have for linking MT activity with experiences of motion is evidence that MT functions as a differentiating correlate of such experiences—we have no evidence whatsoever that it constitutes a total correlate of experiences of motion. Moreover, the thought that MT activity might qualify as a total correlate of such experiences has little to recommend it. One wouldn't expect a slice of MT that had been put in a bottle to generate visual experience, no matter how much current might be run through it. In order to generate visual experience, MT activity must be suitably integrated with non-differentiating activity. Of course, none of the foregoing establishes that holistic approaches to consciousness are more plausible than atomistic approaches. My aim in this section has been merely to demonstrate that what we know about the role played by local neural activity in the generation of consciousness does not provide any support for atomism.

A second argument for atomism appeals to the fact that the components of a stream of consciousness will typically possess distinctive causal profiles. A pain in one's tooth will cause one to visit the dentist, a feeling of fatigue will cause one to take to one's bed, and a thirst for beer will cause one to inspect the contents of one's fridge. The atomist might argue that these commonplace observations indicate that the components of a subject's phenomenal field are atoms, for states with distinctive causal powers must be conceived of as having some kind of robust independence.

We do need to account for the distinctive causal powers of the various components of an overall phenomenal field but we need not embrace atomism in order to do so. States can possess distinctive causal powers without having independent existence; indeed, they can possess distinctive causal powers

without having the *capacity* for independent existence. The weight of an object provides it with one fund of causal powers and its colour provides it with another, but an object's weight and its colour are not independent existents that must be fused in some way. Various features 'enter' consciousness in one fell-swoop, but having entered consciousness they bring with them different causal powers.

It might be useful here to consider a related debate concerning the structure of belief. Common sense (not to mention much philosophy) takes an 'atomistic' approach to the structure of belief, holding that a subject's total belief state is built up out of individual beliefs—the belief that 2+2=4, that Timbuktu is in Mali, that scrambled eggs are best made with a little bit of milk, and so on. On this conception of things, the subject's total belief state is, so to speak, secondary to their individual beliefs. But some have argued that we should instead think of subjects as having only a single state of belief (Lewis 1994; Stalnaker 1984). This state would be global in that its content would include everything that the subject believes. Does this 'top-down' conception of belief imply that the subject's 'sentence-sized' beliefs cannot have distinctive causal powers? I don't think so. Even if I believe that Timbuktu is in Mali and that scrambled eggs are best made with milk only in virtue of being in a single, global belief state which in some way includes each of these two beliefs as elements, it is nonetheless true that my vacation-planning behaviour implicates only the former state whereas my breakfast-making behaviour implicates only the latter state. The causal articulation of one's belief set does not require that that set be built up out of individual, fine-grained beliefs. Similarly, a subject's total phenomenal state is causally articulated, but we needn't think of this causal articulation in atomistic terms.

A final argument for the atomistic approach concerns the modal independence that various phenomenal states have from each other. Take a total phenomenal state (e_1). We can imagine another total phenomenal state (e_2) that is identical to e_1 apart for the fact that the subject lacks one of e_1's components—say, a pain in the left leg (e_3). In other words, we can think of e_2 as e_1 'minus' e_3. But if e_2 qualifies as an 'atom' of experience when it occurs on its own as a total phenomenal state, should it not also qualify as an atom of experience when it occurs in the context of e_1?

I think not. Begin by noting a rather odd implication of the argument. Let us assume that the argument is sound, and that e_2 does indeed qualify as an atom of consciousness when it occurs in the context of e_1. Now, what should we say about the status of e_3—the subject's experience of a pain in the left leg? Although e_2 might have the potential to constitute a total phenomenal state

in its own right, it is highly doubtful whether the same can be said of e_3. Arguably, the experience of a pain in one's left leg can occur only in the context of certain other experiences (such as an overall sense of one's body). It seems to follow from this that e_3 isn't a phenomenal atom, for the argument appears to presuppose that a state qualifies as a phenomenal atom *only if* it has the potential to constitute a total phenomenal state. But if e_3 isn't a phenomenal atom then it's not clear just how we are to think of it. It looks as though the atomist will need to distinguish between two kinds of phenomenal states: phenomenal atoms (such as e_2) and states (such as e_3) that are not phenomenal atoms despite providing a distinctive contribution to the subject's overall experiential state. This position may not be incoherent, but it does not strike me as particularly attractive.

Although this objection suggests that something goes wrong with the argument from independence it doesn't identify where the argument goes wrong. The root of the trouble, however, is not difficult to locate: the fact that a state has the *potential* to constitute an independent unit of consciousness does not ensure that it *is* an independent unit of consciousness. We can see this by returning again to Tye's analogy between experiences on the one hand and statues and clouds on the other (see §2.1). As Tye points out, although certain undetached parts of statues and clouds would constitute statues and clouds in their own right were they suitably detached, as undetached parts they do not. The point generalizes beyond objects like statues and clouds to include events. Consider a football match involving five minutes of stoppage time. Call this event F_1. Now, had the referee blown the whistle at the end of the fourth of those five minutes of stoppage time, the game would have had only four minutes of stoppage time. Call this counterfactual game F_2. Should we hold that because F_2 would have qualified as an 'atom of football' when it occurs on its own then it also qualifies as an 'atom of football' when it occurs within the context of F_1? Surely not. When it comes to understanding football matches—not to mention statues and clouds—there is clearly a great deal to be said for starting with the whole rather than its parts.

Although the argument from independence fails, it does highlight an important feature of consciousness that any account must accommodate: tokens of a single fine-grained phenomenal state type can occur within the context of various total phenomenal state types. (At least, this is true for many fine-grained phenomenal state types; there may be content-based constraints on the types of conscious states that can co-occur within a total phenomenal state (Dainton 2006). However, the essential point is that this fact can be accommodated without embracing atomism.

10.6 Towards a plausible holism

The lesson to be learnt from the preceding sections is clear: our approach to consciousness should be holistic. Rather than begin with atomic states of consciousness that must be 'glued together' to form total phenomenal states we should regard total states as the basic units of consciousness. But what exactly might a plausible version of holism look like? My aim here is not to provide a 'worked-up' model of consciousness, but to develop a metaphor that might be of some use in guiding the construction of such a model.

The metaphor in question likens consciousness to a quilt. Just as multiple squares of cloth are patched together to form a quilt, so too multiple coalitions of content are woven together to form an overarching state of consciousness. One can identify particular fine-grained states of consciousness within this overarching state—just as one can pick out particular squares within a quilt—but these states are not independently conscious. The elements of a subject's total phenomenal state do not 'enter' consciousness as independent units but only *en masse*.

The key contrast between the quilted model and various versions of atomism concerns their respective conceptions of phenomenal unity. In order to account for the unity of consciousness the atomist must posit some kind of mechanism which functions to bind the atoms of consciousness together into phenomenal wholes. The advocate of the quilted model need posit no such mechanism, for the quilted theorist holds that the components of a total state of consciousness come into being *as* the components of that total state. One need no more bind the various components of consciousness together than one need bind the various organs of a body together—they come into being as unified with each other.

Within the contemporary landscape the account that is most congenial to the quilted conception of consciousness is Tononi's dynamic core model, according to which the neural basis of consciousness involves reciprocal interactions between thalamic and cortical processing ('thalamico-cortical loops'). The dynamic core is grounded in sub-cortical systems and reaches out to include various domain-specific processing nodes within its sweep. The very nature of this process ensures that any features that are made conscious are made conscious together—as the components of a single phenomenal state. The thalamico-cortical loops are integrated with each other, such that we should more properly speak of a *single* thalamico-cortical loop rather than multiple loops. Perhaps it should come as no great surprise that Tononi's account appears to fit most closely with the ideas that I have outlined here, for he is one of the few

authors in the recent literature to have taken the unity of consciousness as a serious constraint on theorizing about consciousness (Tononi 2004, 2007; Tononi & Edelman 1998).[13]

Although the dynamic core may have a stable sub-cortical ground it does not have a unitary anatomical location, for the location of the thalamico-cortical loops may vary between subjects and indeed even within subjects across times. Nor does this account imply that we can identify consciousness with some kind of 'box' in an information-processing flow-chart. The architecture of cognition is relatively stable, but that of consciousness is highly labile and changes on a time-scale that is measured in hundreds of milliseconds. Certain phenomenal features such as one's sense of bodily presence and background mood might have an abiding presence in consciousness; others flutter in and out of consciousness depending on perceptual input and one's attentional focus. Further, the boxes of information-processing psychology typically have a restricted domain, whereas the dynamic core can bind an extremely heterogeneous range of content together. Rather than attempting to locate consciousness within the box-and-arrow diagrams beloved by cognitive neuroscience, it might be better to think of consciousness as involving a dynamic unity that is *superimposed* on the relatively static structure of thought and perception.

Does the dynamic core conception of consciousness qualify as a version of the 'imperial' approach to consciousness in the sense that I outlined in Chapter 5? It might appear to, for one might be tempted to think of the 'core' as some kind of 'seat of consciousness': a location within functional space towards which all content flows and from which all control emanates. But to conceive of the core in these terms would be to misunderstand its role, which is merely to generate a total phenomenal state. In this respect, the dynamic core model is perfectly consistent with what I called the federal conception of consciousness, according to which the various components of the stream of consciousness are contained within domain-specific circuits. The dynamic core is in the business of ensuring that contents of various kinds are fused together into a single phenomenal state, but it is not in the business of ensuring that each of these contents will be available to the same range of consuming systems. Nor, for that matter, is it in the business of solving the binding problem or ensuring that the contents of consciousness are consistent with each other. Although the dynamic core might

[13] Also congenial is Hurley's suggestion that consciousness involves a dynamic singularity in the field of causal flows: a tangle of multiple feedback loops of varying orbits (1998, 2003). However, Hurley weds her dynamical approach with a commitment to 'vehicle externalism', the thought that the vehicles of consciousness loop out into the world. Although not opposed to externalism in principle, I am not particularly enthusiastic about it.

play *a* role in ensuring that these kinds of unities are typically present in consciousness, these functions are not themselves the responsibility of the dynamic core.

The quilted model generates no specific methodological prescriptions but it does provide a very general framework for approaching the study of consciousness. Let me bring this section to a close by reflecting on one the questions posed by this framework (Bayne 2007a; Hohwy 2009).

Most investigations into the neural basis of consciousness employ what we might call a *content-based methodology*. Studies designed on this basis hold the presence of consciousness itself fixed and manipulate only the contents of consciousness. The binocular rivalry studies of Logothetis and colleagues that I mentioned in Chapter 5 provide an influential example of this approach.[14] In these studies Logothetis and colleagues trained monkeys to 'report' their visual experiences by pressing different levers, and then recorded from cells in the visual cortex whilst the monkeys were experiencing binocular rivalry. A more recent example of the content-based methodology is provided by a series of studies conducted by Rees and colleagues involving patients with perceptual extinction. The experimenters contrasted the neural activity seen when the visual stimulus was 'extinguished' (that is, not among the contents of the patient's consciousness) with that which occurred when the patient *was* conscious of it.[15]

This content-based approach is in a good position to uncover the differentiating correlates of consciousness, but it is rather more difficult to see how it might enable us to identify the *non-differentiating* correlates of consciousness. This wouldn't be so bad if one thought of non-differentiating activity as merely enabling, but we have found reason to reject that view. Differentiating activity represents only half of the story as far as the neural underpinnings of consciousness are concerned.

In order to identify the non-differentiating correlates of consciousness we need to employ the *creature-based methodology*. Studies conducted in accord with this rubric attempt to screen off the influence of the contents of consciousness and look instead for the domain-general neural activity that correlates with the presence of consciousness as such. An example of this approach is provided by the work of Alkire and colleagues, who have used anaesthesia-induced loss of consciousness to identify the mechanisms responsible for the transition between consciousness and unconsciousness.[16] Other examples of the creature-

[14] See Logothetis (1998); Sheinberg & Logothetis (1997); see also Tong et al. (1998).
[15] See Rees (2001); Rees et al. (2002); see also Driver & Mattingley (1998) and Sarri et al (2006).
[16] See Alkire et al. (2000); Alkire & Miller (2005); White & Alkire (2003).

based methodology are provided by studies that attempt to identify the neural factors which account for the loss of consciousness in coma and the persistent vegetative state and its re-acquisition in the minimally conscious state (see §6.1).[17] Such studies aim to uncover the domain-general, non-differentiating correlates of consciousness rather than their content-specific counterparts.

It is clear that there should be a place for both content-based and creature-based methodologies within the science of consciousness: the former are needed in order to identify the differentiating correlates of consciousness, whereas the latter are needed in order to identify the non-differentiating correlates of consciousness. What is less clear is whether the independent deployment of these two approaches will provide us with a full account of the neural basis of consciousness. I have my doubts. What we really need to know is how differentiating neural activity *interacts* with non-differentiating activity so as to generate the subject's total phenomenal state. It is possible that such a picture will emerge as a result of the independent execution of content-based and creature-based studies, but it is also possible that we will need to develop new methodologies—methodologies that involve the conjoint manipulation of both 'state consciousness' and 'creature consciousness'—in order to determine just how the quilt of consciousness is knit together from the various fragments of content distributed throughout the brain. But just what such a methodology might look like is a question that must be left for another occasion.

10.7 Conclusion

Theory-building in consciousness studies has generally paid scant attention to the unity of consciousness. This is unfortunate, for the unity of consciousness provides us with important and much-needed constraints on accounts of consciousness. I began in §10.2 with a constraint based on object-unity. According to the binding constraint, perceptual features cannot enter consciousness without first being bound together to form percepts of unified objects. Although few accounts of consciousness fall foul of the binding constraint, it is at odds with the kind of radical atomism endorsed by Zeki and Bartels (1999). In §10.3 I provided a brief overview of the vast literature on inter-modal integration, and argued that this body of work puts further pressure

[17] See Laureys et al. (2000); Schiff (2004).

on atomism, for it indicates that the senses do not operate as a 'mass of tributaries running in parallel', as some atomists claim (O'Brien & Opie 1998), but as a network of highly entangled and inter-dependent channels.

But the real pressure on atomism derives from the fact that conscious states occur as the components of a single phenomenal field. How might atomists account for the phenomenal unity of consciousness? They have two options: they can either deny that consciousness is unified, or they can posit a mechanism that might be responsible for ensuring that the atoms of consciousness are brought together in the form of phenomenal wholes. Neither of these options has much to recommend it. The first option flies in the face of introspection, whereas the second generates an unwanted prediction: if the stream of consciousness involves the activity of mechanisms of 'phenomenal binding', then why are there no syndromes in which these mechanisms have broken down leaving the subject in question with phenomenal atoms? And even if—contrary to what I have argued in previous chapters—there are syndromes in which phenomenal unity is lost, the manner of its breakdown surely provides little encouragement for atomists.

The upshot, I suggested, is that the structure of consciousness is fundamentally holistic: there are no mechanisms responsible for phenomenal binding because the unity of consciousness is ensured by the very mechanisms that generate consciousness in the first place. Drawing on Tononi's 'dynamic core' model of consciousness, I suggested that we should think of the mechanisms of consciousness as producing a quilt of consciousness. This quilt is generated by sub-cortical systems that reach out into various cortical nodes and bind distributed fragments of content into a single, multifaceted phenomenal state. Although the various phenomenal states that can be distinguished within this quilt should not be thought of as phenomenal atoms, they do qualify as experiences in their own right, for they may possess distinct causal profiles and a certain degree of functional autonomy.

I brought this chapter to a close by examining one of the many methodological questions that it raises. Investigations into the neural basis of consciousness currently take one of two forms. Some studies ignore creature consciousness and focus on the differentiating correlates of certain kinds of fine-grained conscious states. Other studies ignore the contents of consciousness and focus on the non-differentiating correlates of consciousness that might be implicated in consciousness as such ('creature consciousness'). Both kinds of studies are important, but it is unclear whether either approach is able to reveal how the differentiating correlates of consciousness *interact* with the non-differentiating correlates. Although an account of this interaction might emerge from integrating the results of content-based studies with those of creature-based

studies, but it is also possible that we may need to develop novel experimental approaches—approaches that in some way manipulate both creature consciousness and the contents of consciousness in tandem—in order to determine just how the dynamic core goes about stitching content that is distributed across a variety of neural circuits into a single quilt of consciousness.