

Conscious states and conscious creatures: Explanation in the scientific study of consciousness

Tim Bayne
University of Oxford and St. Catherine's College
Manor Road
Oxford OX1 3UJ
United Kingdom
tim.bayne@gmail.com

This paper is forthcoming in *Philosophical Perspectives* (2007). Please consult the published version for purposes of quotation.

1. Introduction

Explanation does not exist in a metaphysical vacuum. Conceptions of the structure of a phenomenon play an important role in guiding attempts to explain it, and erroneous conceptions of a phenomenon may direct investigation in misleading directions. I believe that there is a case to be made for thinking that much work on the neural underpinnings of consciousness—what is often called the neural correlates of consciousness—is driven by an erroneous conception of the structure of consciousness. The aim of this paper is lay bare some connections between the explanation of consciousness and the structure of consciousness, and to argue for a conception of the structure of consciousness that is more adequate than that which currently drives much research into the neural correlates of consciousness.

The form of consciousness with which I am interested is phenomenal consciousness. We can think of phenomenal consciousness in two ways: in terms of phenomenal states, or in terms of phenomenally conscious creatures. Phenomenal states are individuated in terms of their phenomenal character—that is, in terms of what it is like to instantiate them. There is something it is like to hear a siren sound, to see a field of daffodils in spring, and to smell freshly brewed coffee. Arguably, there is also something it is like to think that $2+2=4$, intend to put the washing out, and wonder when house prices will fall. Phenomenally conscious creatures are creatures that there is something it is like to be.

Some entities are phenomenally conscious; others are not. A theory of creature consciousness will tell us what it is that distinguishes the former from the latter.¹

It is tempting to take the contrast between state consciousness and creature consciousness as a contrast between the type of entity that is conscious: state consciousness is had by mental states whereas creature consciousness is had by creatures. But in fact, I think that this is not the best way to think of the state/creature contrast, for both state consciousness and creature consciousness are properties of creatures. Instead, the contrast between creature consciousness and state consciousness is best thought of as a distinction between a determinable and its determinates: conscious states are determinates of creature consciousness. We might compare the relationship between creature consciousness and state consciousness with the contrast between being coloured and being red, or the contrast between being square and being shaped.

The contrast between creature consciousness and state consciousness presents those studying the neural basis of consciousness with a question: which of these two types of properties should they focus on? Should they look for the neural correlates of creature consciousness? Should they look for the neural correlates of particular states of consciousness? Or should they attempt to find the correlates of creature consciousness and state consciousness in tandem? I examine these questions against the backdrop of debates concerning the structure of consciousness. Some theorists argue that consciousness has a building block structure, others argue that it has a field-like structure. The building block model of consciousness seems to suggest that the study of consciousness should focus on the explanation of phenomenal states, whereas the field-based model would seem to encourage a creature-based approach to the study of consciousness. I present some broad reasons to favour a field-based conception of the structure of consciousness, and conclude with some reflections on how the science of consciousness ought to proceed in light of that result.

2. States, creatures and the neural correlates of consciousness

Studies of the neural mechanisms underpinning consciousness can be grouped into two broad classes: state-based studies and creature-based studies. State-based studies focus on the neural mechanisms underpinning particular (fine-grained) phenomenal states. An

¹ My use of 'creature consciousness' differs from another use that currently has some currency. A number of theorists use 'creature consciousness' for wakefulness (e.g. Carruthers 2005; Rosenthal 2005). Being awake is certainly a respectable property, but it is not the one in which I am interested. Dreamers are—or at least can be—phenomenally conscious without being awake.

early (and influential) example of the state-based strategy can be found in the binocular rivalry studies of Logothetis and colleagues (Logothetis 1998; Sheinberg & Logothetis 1997; see also Tong et al. 1998). When the left and right eyes are presented with different images—say, a sunflower and a house—one typically sees either the house or the sunflower, but not both at the same time, even though both stimuli are processed to a high level. Logothetis and colleagues trained monkeys to ‘report’ their visual experiences by pressing different levers, and then recorded from cells in their visual cortex whilst the monkeys were experiencing binocular rivalry. They found that changes in visual experience were most strongly correlated with activity in inferior temporal cortex, and only weakly associated with activity in primary visual cortex.

A more recent example of the state-based approach can be found in the perceptual extinction studies of Rees and colleagues (Rees 2001; Rees et al 2001; see also Driver & Mattingley 1998). Perceptual extinction is a phenomenon in which patients ignore stimuli presented in one visual hemi-field (typically the left) when presented with stimuli in both visual hemi-fields. The stimuli presented in the patient’s good hemifield capture attention, and lead the patient to ignore the stimuli in their bad hemifield. Rees and colleagues found that the difference between extinction and its absence correlated with activity in frontal (and parietal) attentional and working memory systems.

Various criticisms can be levelled against the state-based approach. One might worry that the neural mechanisms underlying one kind of phenomenal state—say, the experience of seeing a house—might not have anything in common with those underlying other kinds of phenomenal states. Proponents of the content-based approach are not unaware of this objection. They admit that the approach involves a certain amount of risk: an account of the neural correlates of one kind of phenomenal state may not generalize to other kinds of phenomenal states. If no such generalization is forthcoming, then various phenomenal states will have to be studied on a case-by-case basis. Still, the proponent of the content-based paradigm might say, there is no principled objection to the state-based approach.

A deeper worry with the approach is that it does not involve the manipulation of consciousness itself. Instead, only particular phenomenal states are used as variables. Studying particular phenomenal states can tell us something about the neural mechanisms needed for the phenomenal states in question, but it cannot tell us anything about consciousness per se. Arguably, we have to turn to another experimental paradigm in order to study consciousness as such: we need to employ creature-based studies.

Creature-based studies manipulate creature consciousness as a critical variable. One set of creature-based studies focuses on the mechanisms responsible for anaesthesia. Although

different anaesthetics function in subtly different ways, they appear to share a common core of activation centred on the thalamus (Angel 1991; Angel 1993). Alkire and colleagues have discovered that anesthetic-induced unconsciousness is associated with a decrease in functional connectivity between thalamus and cortex; they refer to anaesthetics as flicking a thalamocortical “switch” (Alkire et al 2000; Alkire & Miller 2005; White and Alkire 2003).

Another creature-based approach to the study of consciousness examines the neural mechanisms underlying the transition between coma (in which the patient is assumed to have no phenomenal states) and the minimally conscious state (in which the patient is thought to be phenomenally conscious, at least at times). The transition from coma to the minimally conscious state is correlated with increased activation in the sub-cortical arousal systems (Graham et al 2005; Laureys et al 2000; Schiff 2004; Schiff & Plum 1999; Schiff & Plum 2000).

A broad-brush picture of the neural correlates of consciousness is as follows. On the one hand, content-based studies suggest that phenomenal states, such as seeing a face, or hearing a melody, depend on the activation of particular cortical regions—what are sometimes called ‘cortical nodes’. At the same time, creature-based studies suggest that sub-cortical systems centred on the thalamus are necessary for the organism to enjoy any phenomenal states at all. What we need, at this juncture, is an account of how the sub-cortical systems and cortical nodes conspire to produce consciousness. And we need to locate the contrast between content-based studies and creature-based studies in an account of the structure of consciousness itself. So, the next item on the agenda is to examine competing conceptions of the structure of consciousness.

3. Building blocks vs. unified fields

Following Searle (2000), we can distinguish two approaches to the structure of consciousness: the building block model and the unified field model.² Building block theorists take phenomenal states, of roughly the kind described above, as the fundamental units of consciousness. On this view, a person’s total phenomenal state is built up out of fine-grained phenomenal states, where a person’s total phenomenal state captures exactly

² Although I employ Searle’s terminology, I don’t endorse all of what Searle says about the contrast between the building block and unified conceptions of consciousness. Searle tends to use these terms both for competing conceptions of the structure of consciousness and for competing conceptions of how the study of consciousness ought to proceed, whereas I am using them solely for accounts of the structure of consciousness.

what it is like to be them at a particular time. My total phenomenal state is constructed out of a visual experience of a chair, an auditory experience of the music that I am currently listening to, a certain emotional experience, bodily sensations of various kinds, and so on.

As the name suggests, unified field theorists hold that consciousness involves a field. A creature has a particular phenomenal state when the relevant content “enters” a phenomenal field; that is, when the creature’s phenomenal field is modulated in a certain way. A creature’s phenomenal field might be rich in content or poor in content. Whatever its make-up, the unified field approach to consciousness regards phenomenal states as derivative on the existence of a field of consciousness.

One might worry that the difference between the state-based and field-based conceptions of consciousness is more apparent than real. Why shouldn’t we identify phenomenal fields with total phenomenal states? I think that there is a sense in which one can translate talk of phenomenal fields into talk of total phenomenal states, but it proves very convenient to posit phenomenal fields in addition to total phenomenal states. Here is why. Suppose that at t1 I am in a total phenomenal state that includes only visual and auditory experiences, while at t2 I not only enjoy visual and auditory experiences but also a certain olfactory experience (say, the smell of freshly roasted coffee). My total phenomenal states at t1 and t2 are different, and it is tempting to capture this difference by saying that my field of consciousness has ‘expanded’, so to speak—it is richer than it was 5 seconds ago. But states cannot expand or contract, for they are individuated in terms of their content. So, rather than think of the stream of consciousness as a succession of total phenomenal states, we might want to think of it as involving the activity of a phenomenal field, the contents of which change through time.

The proponent of the field-based approach to consciousness might be tempted to hold that there are no phenomenal states, there are only phenomenal fields and their contents.³ This would be a mistake. There is no reason for the proponent of the field-based model to deny that there is such a thing as the experience of seeing something as red, and that different subjects of experience can have such an experience. The point that the field-based theorist will emphasize is that experiencing something as red should not be regarded as a primitive state. It might be useful here to consider a parallel with theories of belief. It is sometimes suggested that, strictly speaking, individuals have a single belief-

³ Although he doesn’t use the ‘phenomenal field’ terminology, Tye’s (2003) approach to the structure of consciousness has certain affinities with the field-based account. However, Tye argues that fine-grained phenomenal states do not exist, which I think is an over-reaction to the building-block model. See Bayne (2005) for an analysis of Tye’s position.

state with rich content rather than lots of beliefs (see Lewis 1994; Stalnaker 1984). But note that one can adopt this view without rejecting all talk of fine-grained belief states. We can refer to the belief that mammals are dogs, even if—strictly speaking—the belief that mammals are dogs is an abstraction from the more complex belief-state that the person enjoys. Similarly, we can speak of the experience of smelling freshly roasted coffee, even if—strictly speaking—such an experience is (almost always) only one component of the person’s overall phenomenal field.

We can put some flesh on the distinction between these two conceptions of the structure of consciousness by asking who might be thought to espouse each of them. I will focus only on neurobiological accounts of consciousness, and will leave philosophical theories of consciousness to one side here.

Searle describes Crick and Koch as paradigmatic building block theorists. There is something to what Searle says, for there is certainly a building block orientation to Crick and Koch early (and influential) “40 Hz” model of consciousness, according to which consciousness is subserved by the mechanisms responsible for binding various perceptual features together to form percepts of unified perceptual objects (‘dynamic object binding’). This model suggests that a person’s overall phenomenal perspective is built up out of units of perceptual objects.

But it is far from obvious that Crick and Koch are dyed-in-the-wool building block theorists. In recent work they describe frameworks which have a unified field orientation—such as Edelman and Tononi’s dynamic core model and Dehaene and Naccache’s global workspace account—as ‘somewhat similar’ to their own (Crick and Koch 2003: 124).⁴ Furthermore, they recognize that domain-general enabling factors are needed for consciousness.

Our main interest is not the enabling factors needed for all forms of consciousness, such as the activity of the ascending reticular systems in the brainstem. Rather, we are interested in the general nature of neural activities that produce each particular aspect of consciousness, such as perceiving the specific colour, shape or movement of an object (Crick & Koch 2003: 119)

⁴ We should note that even in their earlier work Crick and Koch were not wholly committed to the building block model. For example, they described content integration as involving a “form of short-term memory and serial attentional mechanism, such that a temporary global unity is imposed on neurons in many different parts of the brain” (Crick and Koch 1990: 1).

The crucial issue here concerns what Crick and Koch mean by an “enabling factor”. On the one hand, they might regard the activity of the ascending reticular systems as merely *causally* necessary for consciousness. On this view, cortical states form the total neural correlates of phenomenal states—the reticular system contributes to consciousness only insofar as it is needed to activate the relevant cortical nodes. Were one to get the cortical nodes into the right state of activation, they would generate the relevant phenomenal states under their own steam, so to speak. On the other hand, Crick and Koch might regard the activity of the ascending reticular systems as *constitutively* necessary for consciousness. On this view of things, the role of the reticular system is not (just) that of activating the subject’s cortical nodes, rather, this system is an integral component of the total neural correlate of the subject’s phenomenal states.⁵ It is not clear which of these two models Crick and Koch endorse, but only the former entails the building block approach. Someone who comes closer to an explicit endorsement of the building block model is Nancy Kanwisher.

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).

Whether Kanwisher really means to be endorsing the building block model depends on how exactly we are to understand the final sentence. Does she mean to suggest that the *total* correlate of the subject’s visual experiences can be located in content-specific cortical nodes, or does she mean to suggest only that those experiences are dependant on

⁵ On a functionalist account, one might say that the supervenience base of a conscious state P is a supervenience base only insofar as it has causal effects elsewhere.

particular cortical nodes? The unified-field theorist need not disagree with the latter claim, but I suspect that Kanwisher has the former thesis in mind.

Perhaps the most explicit endorsement of the building block model, at least within the cognitive neurosciences, comes from Zeki and Bartels.

Activity at each stage or node of a processing–perceptual system has a conscious correlate. Binding cellular activity at different nodes is therefore not a process preceding or even facilitating conscious experience, but rather bringing different conscious experiences together. (Bartels & Zeki 1998: 2330; see also Zeki 2007).

Zeki and Bartels call these conscious experiences ‘micro-consciousnesses’. As Zeki and Bartels conceive of them, micro-consciousnesses are independent conscious states—phenomenal building blocks.

What about unified field models of consciousness? Perhaps the model of consciousness most naturally deserving of the ‘phenomenal field’ title is Tononi’s information integration theory of consciousness (Tononi 2004; Tononi 2007; see also Tononi & Edelman 1998). According to this model, information is conscious in virtue of being caught up in a highly integrated neural process. Tononi suggests that although the neural substrate of this complex will be dynamic rather than fixed, it will involve thalamocortical circuits. One could also make a case for thinking of global workspace models of consciousness in unified field terms. According to such models, content becomes conscious in virtue of entering a global workspace—a sort of domain-general working memory from which content is available to a range of high-level consuming systems (Baars 1988; Dehaene & Naccache 2001). The global workspace model of consciousness is suggestive of the unified field approach insofar as it regards various contents as being made conscious in virtue of the operations of a single system—the global workspace.

4. Structure and methodology

How does the distinction between the building block and unified field models of the structure of consciousness relate to the distinction between state-based and creature-based approaches to the study of consciousness?

On the face of things, the building block model seems to recommend a state-based approach to the study of consciousness, according to which the study of consciousness ought to focus on phenomenal states. The building block model suggests that there is little point in studying the mechanisms of creature consciousness, for on the building block model creature consciousness is derivative on state consciousness. The unified field account, by contrast, would seem to recommend a creature-based approach to the study

of consciousness, for on the field-based model a creature enjoys the particular phenomenal states that it does only because of the ways in which its phenomenal field is modulated. We might put the contrast as follows: the phenomenal field theorist sees the determinable of creature consciousness as having explanatory priority, whereas the building block theorist sees the determinables of creature consciousness—that is, fine-grained phenomenal states—as having explanatory priority.

However, certain important complications to the foregoing must be noted. Firstly, state-based studies are not equipped to reveal the building blocks of consciousness. The reason for this is that the state-based approach does not allow one to determine the *total* neural correlate of a particular phenomenal state, instead, it allows one to identify only that neural state that is correlated with the particular phenomenal state in question *in an already conscious subject*. State-based studies show only that activation of the fusiform face area generates an experience of seeing a face in the context of an already conscious creature, and so we cannot employ these studies to identify the *total* neural correlate of the visual experience of a face. At best, such experiments can reveal those neural systems that are correlated with the contents of the phenomenal states in question.⁶

We might say that the state-based approach allows one to determine the neural correlates of the *contents* (or phenomenal character) of consciousness, but even this way of putting things is tendentious. Suppose that we contrast two situations: in S1 the subject sees a house, in S2 the subject sees a face. Should we identify the neural states that are distinctive of S1 and S2 with the correlates of the contents <seeing a house> and <seeing a face> respectively? No, for S1 and S2 are both visual experiences, and, as such, we might expect them to share a common neural correlate in virtue of their visual character.

What about the link between the phenomenal field model and the creature-based methodology? Here too words of caution must be sounded. It is tempting to assume that the mechanisms responsible for generating creature consciousness are identical to the mechanisms responsible for the phenomenal field, but matters might be slightly more complicated than this. If a creature can have multiple phenomenal fields at a time, then there will be more to the explanation of a phenomenal field than the explanation of creature consciousness. And there does not seem to be anything incoherent in the notion

⁶ Some theorists use the term 'core-NCC' to refer to the state that is uncovered by these experiments (Block 2005; Chalmers 2000). I avoid this term because it downplays the importance of the neural mechanisms subserving creature consciousness. Arguably, such mechanisms are more deserving of the title 'core-NCCs' than are the cortical nodes that are associated with particular (fine-grained) phenomenal contents.

of a single organism having multiple phenomenal fields at a time (see the discussion of the split-brain below). But leaving the possibility of multiple phenomenal fields to one side, the phenomenal field conception of the structure of consciousness does indeed recommend the creature-based approach to the study of consciousness. But this is not to say that the unified field theorist must condemn state-based approaches as impotent—far from it. In section 6 I will sketch a way in which state-based methodologies might allow us to uncover the mechanisms of consciousness, even if the phenomenal field conception of consciousness is correct. But first I turn to the debate between the building block and unified field models of consciousness.

5. In defense of the phenomenal field

The debate between the state-based and field-based models of consciousness cannot be settled on conceptual grounds. The mere analysis of the concept of consciousness will not pronounce on this question. Each model is coherent, and the debate between them is an empirical one. Nevertheless, it seems to me that there are certain very broad reasons—reasons that are already available to us—to think that consciousness has a field-based structure.

5.1 Phenomenal unity

Human beings can, and routinely do, enjoy a fairly rich phenomenal life. Not only do we enjoy experiential content across a wide range of perceptual modalities, we also enjoy a wide array of experiential content within perceptual modalities. Now, although it is something of an open question just how much phenomenal content can be loaded into consciousness at any one point in time, there is good reason to think that quite a lot of content can simultaneously enter consciousness. One can be simultaneously conscious in a variety of perceptual modalities, whilst also enjoying a sense of embodiment, a background sense of affective tonality, and a stream of inner thought and reflection.

These various conscious states do not occur as phenomenal atoms. Rather, they are experienced together, as components of single total phenomenal state. At any one point in time, one enjoys a single phenomenal perspective that subsumes whatever particular phenomenal states one has at the time in question (Bayne & Chalmers 2003; Dainton 2000).

This conception of the unity of consciousness has two grounds. Firstly, it is suggested by introspection: we typically experience ourselves as having a single all-encompassing

phenomenal state.⁷ Secondly, it is suggested by the integration of thought and action. Practical and theoretical reasoning involves integrating the contents of perceptual experience, affective experience, bodily sensations and various states of cognitive consciousness.

In order to account for this unity, the building block theorist must appeal to mechanisms whose job it is to stitch the building blocks of consciousness together so as to form a total phenomenal state. We might call this process of stitching together the subject's phenomenal building blocks 'phenomenal binding'. Zeki and Bartels appear to invoke such a process when they refer to binding cellular activity as "bringing different conscious experiences together".

It is important to distinguish phenomenal binding from object binding. Most discussions of binding are concerned with the question of how representations of perceptual features (colour, shape, motion, etc) are put together to form representations of unified perceptual objects. Everyone needs to posit mechanisms of object binding, but only building block theorists need to posit mechanisms of phenomenal binding.

There are two problems with the phenomenal binding proposal. The first problem concerns timing. The intuitive picture painted by the building block theorist is that the building blocks of consciousness are first constructed, and only then are they bound together to form total phenomenal states. But this would appear to entail that there is a period during which the building blocks exist as unbound, and this in turn would be at odds with introspection, for we clearly do not experience a rapid series of disjoint and then unified experiences, only to have them replaced by another series of disjoint and then unified experiences.

This objection is not decisive, for the building block theorist need hold only that there are two types of mechanisms at play in the construction of the subject's total phenomenal field: one set of mechanisms generates phenomenal states, and another set of mechanisms binds the resulting phenomenal states together.

⁷ At least, this is how things strike me. Not everyone concurs. Some theorists deny consciousness seems to be unified, even in the normal case. Other theorists grant that consciousness seems to be unified, but they insist that this appearance is an illusion. According to these folk, we experience ourselves as having unified phenomenal perspectives because those conscious states of which we are introspectively conscious will be necessarily experienced as unified. But, so the objection goes, it does not follow that all conscious states are unified, it follows only that those conscious states of which we are introspectively conscious are unified. Perhaps there is something to this line of thought. But even if it is right, the point remains the consciousness is typically unified.

A more troubling objection concerns the possibility that the mechanisms of phenomenal binding might malfunction such that the subject in question would enjoy unbound phenomenal states. The mechanisms responsible for generating conscious states could continue to function normally, but the states produced by them would be phenomenal atoms—they would not be subsumed by any total phenomenal state.

Is there any reason to think that such phenomenal disunity occurs? I do not know of any. We simply do not find subjects with ‘orphaned’ phenomenal states. Of course, there are various pathologies in which *some* form of the unity of consciousness is lost. For example, patients with apperceptive agnosia fail to experience the various parts of perceptual objects as the parts of single perceptual objects, and so in some sense they have disunified experiences. But this kind of disunity is not the kind of disunity that would result from a break-down of phenomenal binding, rather, it is the kind of disunity that would result from a break-down of object binding. We should not think of the patient’s experience of each of the parts of the cup as disunified from each other; rather, the problem that the patient has is one of integrating the contents of experience. The patient enjoys a total phenomenal state—a state that includes various representations of the cup—it’s just that this state doesn’t include a representation of the cup as a single, spatially integrated entity.

Arguably, the best case to be made for phenomenal disunity in human beings concerns epileptic patients who have undergone section of the corpus callosum—so-called split-brain patients. Under carefully controlled laboratory conditions, split-brain patients sometimes behave as though they have two conscious fields (streams of consciousness), one in each hemisphere (see e.g. Marks 1981; Zaidel et al 2003). For example, the split-brain patient might be shown the word ‘key-ring’, so that ‘key’ falls within the patient’s left visual field and ‘ring’ falls within the patient’s right visual field. The contralateral structure of the visual system ensures that stimuli projected to the left visual field are processed in the right hemisphere and vice-versa. When asked to report what she sees the patient in the key-ring experiment will typically say only that she sees the word ‘ring’, yet, with her left hand the patient may select a picture of a key and ignore pictures of both a ring and a key-ring. According to the two-streams model of the split-brain syndrome, this disunified behaviour is a manifestation of the fact that the split-brain patient has two streams of consciousness, one in each of her two hemispheres.

Although it is the received view, there is good reason to regard the two-stream account with some suspicion. For one thing, split-brain patients do not display this kind of behavioural disunity outside of carefully controlled laboratory settings. In everyday life, they behave in the integrated way in which someone with a single stream of

consciousness would (Zaidel 1994).⁸ Furthermore, it is possible to explain the disunified behaviour of the split-brain patient without supposing that she or he has a disunified consciousness. In order to do this, we need only suppose that the patient has a single phenomenal field whose contents “switch” between the two hemispheres in accordance with shifts in hemispheric activation (Levy 1977; 1990; Levy & Trevarthen 1976). According to this model, the split-brain operation does not split the patient’s stream of consciousness into two, but merely limits the degree to which the patient’s phenomenal field can simultaneously draw on the neural activity of both hemispheres. In effect, the model represents the split-brain patient as operating with only half of a phenomenal mind at a time. The effects of this in everyday life are minimal because activation moves back and forth between hemispheres as and when it is needed. The switch model derives some support from the fact that the sub-cortical systems concerned with modulating attention and arousal are not seriously damaged by the commissurotomy operation.

But the switch model is controversial and may well be false. Suppose that split-brain patients do have a disunified consciousness: would this provide support for the building block model? I think not. The streams of consciousness that, according to the two-streams model, split-brain patients enjoy hardly look like phenomenal building blocks. There is some elasticity within the building block approach as to how ‘big’ the building blocks of consciousness are, but it is implausible to suppose that the typical phenomenal field is composed of only two building blocks, one grounded in each hemisphere. Even on the two-streams account of the split-brain we should not think of the split-brain syndrome as revealing the pre-existing structure of consciousness.

Let us leave the split-brain syndrome to one side and return to the general issue of phenomenal unity. With the possible exception of split-brain patients, human beings enjoy unified streams of consciousness. This datum needs to be explained. The unified field model has a simple explanation: consciousness has a field-like structure, and any content within this field is automatically unified. More specifically, the very mechanisms

⁸ In response to this fact, some theorists have suggested that split-brain patients have two streams of consciousness only under laboratory conditions, and that in normal contexts they enjoy a single total phenomenal state (see e.g. Tye 2003). This account appears to do justice to both the everyday behavioral unity of the split-brain patient and the behavioural disunity that he or she exhibits under laboratory conditions, but it has problems of its own. The proponent of this account needs to explain how the onset experimental conditions might split a previously unified stream of consciousness into two streams, and how a return to normal cognitive environments reintegrates the patient’s two streams of consciousness.

that generate consciousness also bind the contents of consciousness together so as to form a single (total) phenomenal state. On this picture, there is no need to bind separate phenomenal states together so as to form a total phenomenal state for, strictly speaking, there are no separate phenomenal states to be bound together. Instead, there is simply the phenomenal field with its various contents.

5.2 Phenomenal holism

A second feature of consciousness that appears to be at odds with the state-based model is its holism. But there are a number of senses in which consciousness might be said to be holistic, and not all forms of phenomenal holism are at odds with the building block model.

Firstly, consciousness might be said to be *semantically* holistic. According to the semantic holist, a particular cortical state is correlated with a particular phenomenal state only because that cortical state realizes a certain role in the functional system in which it is embedded. The semantic atomist, by contrast, claims that a neural state realizes a phenomenal state because of some direct relation between it and the property represented by the phenomenal state.

There is much to be said on behalf of semantic holism. Of course, many cells have very specialized receptive fields: cells in the fusiform face area respond selectively to images of faces (Allison et al. 1999; Ishai et al. 1999; Kanwisher et al. 1997; McCarthy 1997); cells in parahippocampal place area respond selectively to images of places (Epstein et al 1999); and cells in medial temporal lobe respond selectively to particular individuals, such as Halle Berry (Quiroga et al. 2005). But none of this tells against the claim that neural representation is holistic. It could be the case that particular cells have the receptive fields that they do only because of their location in the functional architecture of the brain.

Does semantic holism pose a problem for the building block model? Not as far as I can see. Even if neural states carry the contents that they do only because of their role in the overall architecture of the mind, it is entirely possible that neural states are able to simultaneously realize distinct phenomenal states under their own steam, so to speak.

A second form of holism concerns the transitions between consciousness and unconsciousness. Such transitions are typically global: they apply to all forms of phenomenal states, rather than to a select sub-set of the subject's experiences. Typically, on emerging from a dreamless sleep one doesn't first regain visual experience, then auditory experience, then affective experience and so on. Consciousness may dawn gradually, but it dawns gradually over the whole. Similarly, when consciousness is lost—

as, again, might occur in sleep or coma—it is typically lost en masse rather than in a piecemeal fashion. We can capture this fact by saying that consciousness is punctate.

The phenomenal field model can explain the punctate nature of consciousness.

Consciousness is lost or regained when the phenomenal field is lost or regained. Can proponents of the building block model also explain why consciousness is punctate? I think that they can, although the explanation here is perhaps not quite as straightforward as it is in the case of the phenomenal field model. The building block theorist might say that consciousness is punctate because all of the subject's phenomenal states enjoy a common set of enabling factors. When these enabling factors go down, so too does an entire range of phenomenal states. Think here of a power station: when it malfunctions all of the electrical appliances in your house go down with it. But the proponent of the building block model will urge that although a range of phenomenal states depend on a common core of 'enabling factors', these states nonetheless have distinct total neural correlates.

The building block model can account for the two forms of holism just examined, it struggles to account for what we might call the *dynamic* holism of consciousness. Consider the following vignette:

You are starting to cross a street, immersed in conversation with a friend, when suddenly a car swerves towards you. You instinctively take evasive action, and only then does a wave of fear, followed by relief, sweep over you.

There is nothing particularly unusual in this scenario. Entry into the stream of consciousness involves a competition between coalitions of contents, each one of which struggles to make its voice heard above that of its fellows (Koch 2004). Those that win the competition—whether by top-down control or stimulus-driven attention—enter the stream of consciousness; the losers hover in the phenomenal fringes or outside of consciousness altogether. The forces behind these dynamic changes in the contents of consciousness are holistic. In order to predict whether a particular cortical node will make a contribution to consciousness one has to know how strongly other cortical nodes are activated. What matters is not absolute but relative strength of activation: a weakly activated node might dominate consciousness if it is more strongly activated than its competitors, and a strongly activated node might be crowded out by nodes that are even more strongly activated. Whether or not content becomes conscious depends on global factors – that is, on how the mechanisms underlying it are related to the rest of what the organism is doing, its attentional focus, and so on.

This picture is one that is not easily reconciled with the building block model, for the building blocks of consciousness are most naturally thought of as operating independently of each other. Certainly this is the view that some building block theorists have taken. According to O'Brien and Opie, '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' (1998: 387). The dynamical holism of consciousness suggests that consciousness is not built up out of a number of distinct consciousness-generating processes running in parallel, but instead involves a single process into which various contents compete to dominate.

In response, the building block theorist might say that the building blocks of consciousness are highly inter-dependent: the activation of one building block depends in complex ways on the activation of others. We could liken the various states to the members of a committee meeting: whether or not one person speaks might be dependent on whether or not other members of the committee have decided to speak; in some committees one might have to be very aggressive in order to get one's voice heard, whilst in other committees one might have an easy job of dominating proceedings. But here we are in danger of collapsing the building block model into a version of the unified field model, for highly-interdependent mechanisms are best understood as the components of a single consciousness-generating mechanism.

5.3 Background states of consciousness

A third source of support for the unified field model derives from reflecting on the nature of background states of consciousness. It is easier to give examples of background states of consciousness than it is to say exactly what they are. The background state of consciousness with which we are most familiar is ordinary wakefulness, typically referred to in the clinical literature as a state of 'clear consciousness'. Here, one is able to reason through problems, exert some control over one's attentional focus, and one enjoys familiar conscious contents. A quite different background state of consciousness occurs in REM dreaming, in which the contents of consciousness derive from internal stimulation rather than environmental input, reasoning and decision-making is impaired, and one enjoys unusual (at least, relative to normal waking consciousness) conscious content. Other background states of consciousness occur in the state of inebriation and under the influence of various forms of hallucinogens. The hypnotic state is plausibly thought of as involving a particular background state of consciousness, as is delirium, which is often referred to in the medical literature as 'clouded consciousness'.

In the clinical and neuroscientific literatures background states of consciousness are typically referred to as “levels of consciousness”: the comatose state and ordinary wakefulness are represented as marking out the two ends of the consciousness spectrum, with general anesthesia, deep sleep, REM sleep, light sleep, and drowsiness positioned between those two extremes (see e.g. Laureys 2005). I think this conception of background states is problematic, for it falsely suggests that one is more conscious when ordinarily awake than (say) when drowsy or in an REM dream state. Rather than conceive of background states as levels or degrees of consciousness, we should think of them as distinctive regions in the space of phenomenal possibility.

As such, background states are best thought of as ways in which the subject’s phenomenal field is modulated. Because of this, background states of consciousness are domain-general. They characterize the subject’s overall phenomenal perspective, rather than some particular subset of the subject’s phenomenal states. One does not begin dreaming in only some modalities—rather, one’s entire stream of consciousness switches from waking consciousness to dreaming (and back again). Similarly, the transition between normal waking consciousness to hypnosis and back again occurs en masse rather than in a piecemeal fashion. Background states of consciousness are properties of creatures, and they impact on the creature’s particular conscious states only derivatively.⁹

The building block model has no obvious account of why background states should be global (domain general). Why shouldn’t background states of consciousness be restricted to certain components of consciousness if, as the building block theorist asserts, the various components of consciousness are subserved by distinct mechanisms? It seems more parsimonious to assume that background states of consciousness are global because the mechanisms responsible for them are intimately related to those which underlie the phenomenal field, than to suppose that they are global because they affect a number of distinct consciousness-generating mechanisms, as per the building block model.

We might note there is some neuroanatomical plausibility to this framework. The transition between waking consciousness and REM consciousness is correlated with changes to the thalamic system. Roughly speaking, the sub-cortical mechanisms

⁹ In certain situations background states of consciousness might be restricted to a sub-set of the subject’s overall phenomenal states. Perhaps the lucid dreamer and the daydreamer both enjoy both normal waking conscious states commingled with dream states. And perhaps, in certain cases, hypnotized subjects are able to preserve an island of intact normal awareness in the midst of an overall state of hypnosis. But such cases are both obscure and unusual; in general, background states of consciousness appear to be global.

responsible for modulating the subject's background state of consciousness are closely related to those that enable the creature to be phenomenally conscious in the first place. These facts provides additional reason to think of the subject as having a single global state of consciousness into which various contents feed rather than a variety of conscious states that must be bound together by some process that is independent of consciousness itself.

6. Objections and Replies

I turn now to some objections to the unified field approach.

6.1 Phenomenal conservatism

One line of objection to the phenomenal field model involves an appeal to phenomenal conservatism. Phenomenal conservatism is the view that the contents of phenomenal consciousness are restricted in scope. The phenomenal conservative holds that only low-level perceptual, sensory and affective states are phenomenally conscious; high-level perception and cognition are excluded from the realm of phenomenal consciousness (Carruthers 2005; Tye 1996). Phenomenal conservatives typically allow that high-level perceptual and cognitive states have a phenomenology in the derivative sense that the images and bodily sensations associated with perception and cognition possess phenomenal character, but they deny that high-level perceptual and cognitive states have their own phenomenal character. The conservative will grant that high-level perceptual and cognitive states are conscious in *some* sense of the term, but he or she insists that these states are not phenomenally conscious.

From the perspective of the conservative phenomenal consciousness is a relatively minor component of consciousness, and it is unlikely that the mechanisms subserving phenomenal consciousness will be constitutively connected to those subserving consciousness as such. It might be over-stating things to say that phenomenal conservatism is *inconsistent* with the phenomenal field model, but the two accounts certainly make for odd bed-fellows. One can see this by noting that on the conservative approach it is at least conceptually possible for a creature to lose all of its phenomenal states without losing consciousness. And if this is possible, then one cannot ground an account of phenomenality in an account of creature consciousness.

But although phenomenal conservatism has its advocates, it seems to me—as it has to others (Dainton 2000; Goldman 1993; Horgan and Tienson 2002; Pitt 2004; Strawson 1994)—that the range of phenomenal consciousness extends far beyond qualitative, affective and low-level perceptual states. Not only does phenomenal consciousness include high-level perceptual states such as recognizing a face, a voice or a melody, it also

includes various ‘cognitive feelings’—such as the tip of the tongue experiences, feelings of knowing, and déjà vu and jamais vu experiences—and, more controversially, purely cognitive states, such as understanding an argument, judging that one ought to attend a meeting, remembering what one did last winter, and desiring that it stop raining. The phenomenology of these states is not exhausted by whatever bodily and affective sensations accompany them, but belongs to these very states themselves. And if phenomenal consciousness is as much a matter of cognition as it is of perception and sensation, then it is not implausible to suppose that the very mechanisms that enable a creature to be conscious at all might be constitutively related to the mechanisms that allow it to enjoy particular states of phenomenal consciousness.¹⁰

6.2 The empty field?

The model of consciousness presented here has two basic components: mechanisms that generate the phenomenal field, and mechanisms that are responsible for particular contents of consciousness. The former appear to be centred on the thalamic system and associated brain-stem systems; the latter are distributed throughout the cortex (although sub-cortical systems might also contribute to certain contents, particularly those related to affect and bodily sensation). To a first approximation, the organism is conscious in a specific way when both the thalamic system and the appropriate cortical node are activated.

Now, suppose that the system responsible for the phenomenal field is active but all of the cortical nodes responsible for the contents of consciousness are dormant. One might argue that in such a scenario the subject would have an empty phenomenal field. But, so the objection continues, this is not a coherent scenario: one cannot be conscious without being in a particular state of consciousness. A state of objectless (or contentless) consciousness is impossible. In order for a determinate to be instantiated, one (or more) of its determinates must also be instantiated. So the phenomenal field cannot be independent of state consciousness, as the field-based model has it, but must be parasitic on it.

Some might take issue with the claim that an empty (contentless, objectless) phenomenal field is impossible, but I am inclined to allow this claim: in order to be phenomenally conscious one must be phenomenally conscious in some way or another. Nonetheless, the objection can be side-stepped. The crucial point to note is that we cannot *identify* creature

¹⁰ I am tempted to defend the stronger claim that the basic form of consciousness is phenomenal consciousness, and that we lack a notion of cognitive consciousness that is truly independent of the notion of phenomenal consciousness. But I won't attempt to argue for that position here.

consciousness with the phenomenal field. There is nothing it is like to be a creature unless its phenomenal field is modulated in a particular way, and a creature that there is nothing it is like to be is not (phenomenally) conscious. So, a full account of creature consciousness will need to involve an account of both the neural systems underlying particular phenomenal states and the phenomenal field. We might think of the total neural correlate of every phenomenal state as having two components: a component that reflects the content of the state in question, and a component that reflects the activity of the subject's phenomenal field. Phenomenal states will have distinct neural correlates in accordance with their content, but they will have a neural correlate in common insofar as they each involve the activity of a single phenomenal field. We might put the model in Kantian terms: a phenomenal field without content is empty, but content that does not occur within a phenomenal field is blind.

6.3 The integration of contents and fields

The picture of consciousness that I have presented suggests that in order for a subject to be in a particular phenomenal state both the system(s) responsible for the creature's phenomenal field and the system(s) responsible for the content (or character) of the phenomenal state must be activated. But does the activation of these two systems *suffice* for generating the target phenomenal state?

There is reason to think not. Research on visual extinction has found robust category-specific activation in ventral visual cortex even for those stimuli of which neglect patients appear not to be conscious (Rees 2001; Rees et al 2002; Driver & Mattingly 1998). On the face of things, this research indicates that the total correlate of a particular phenomenal state cannot be equated with the mere *conjunction* of the activation of the organism's phenomenal field and the relevant cortical node. Something more seems to be required. What might that more be?¹¹

One possibility is that category-specific activation ('cortical nodes') generate consciousness only when their contents are made available to—or perhaps even contained within—a global workspace (Dehaene & Naccache 2001; Koch 2004; Passingham & Lau 2006). Roughly speaking, the function of this system would be to make the contents of

¹¹ One could challenge the assumption that these patients are not conscious of the stimuli that they neglect. Block (2007) suggests that neglected stimuli might be conscious but unreportable. Although I think that Block's proposal should be regarded as a live option, I will proceed on the assumption that the subject is not phenomenally conscious of neglected stimuli.

consciousness jointly available to the subject's high-level consuming systems (executive planning, problem solving, memory consolidation, verbal report, and so on).

I'm not much attracted to this position, for it seems to me that consciousness—*phenomenal* consciousness—demands little in the way of high-level cognitive availability. Patients with pre-frontal damage are unable to generate coherent and effective action plans, but there is no reason to deny that they are perceptually (or affectively) conscious.

Furthermore, hydranencephalic children, who are born with, at best, only the remnants of a cerebral cortex, show signs of being phenomenally conscious, suggesting that a pre-frontal cortex is not needed for consciousness (Merker 2006).¹²

A better response to the objection, it seems to me, is to hold that particular phenomenal states require the *integration* of cortical node activity and sub-cortical systems. Rather than think in terms of two neural correlates—one of which reflects the content of the state and one of which reflects the activity of the phenomenal field—it is better to think in terms of a single neural correlate: a cortico-thalamic loop. Roughly speaking, the cortical component of this loop might differentiate this phenomenal state from the creature's other phenomenal states, while the thalamic component of this loop might be had by each of the creature's phenomenal states. On this model, being integrated into appropriate corticothalamic activity is both necessary and sufficient for cortical activity to be reflected in the contents of phenomenal consciousness.

These reflections point to one way in which state-based paradigms for the study of consciousness might be employed with a unified-field conception of the structure of consciousness. What differentiates those mental states that can be conscious from those that are unable to enter consciousness? And of those states that are potentially conscious, which distinguishes those that actually enter consciousness from those that are only potentially conscious? The answer to this question, I suspect, concerns the degree to which the states in question can—and do—interact with the mechanisms underlying the creature's phenomenal field. Whether or not content enters consciousness will depend on the relations between those nodes that subserve it and those systems that subserve the creature's phenomenal field. By looking at which cortical nodes support consciousness and which do not we might begin to get a fix on those systems that subserve the phenomenal field. The phenomenal field theorist will not be interested in the intrinsic properties of those cortical nodes that support conscious content, but will rather be interested in the relations that distinguish those cortical nodes that are capable of

¹² Could we regard the frontal-parietal network as a component of the phenomenal field? No, for despite damage to this network patients with neglect are obviously phenomenally conscious.

supporting consciousness from those cortical nodes that are not capable of supporting consciousness.

7. Conclusion

Despite significant progress, the scientific study of consciousness is still very much in its infancy; not surprisingly, it includes competing research agendas. When viewed from afar, two general approaches can be discerned: some studies focus on the mechanisms underlying the contents of consciousness, others focus on the mechanisms underlying creature consciousness. Running in parallel to this methodological division is a debate about the structure of consciousness. Is a person's overall phenomenal perspective built up out of phenomenal building blocks, or does it involve a single phenomenal field into which various contents feed?

I have argued that there is much to recommend a field-based conception of consciousness. The mechanisms subserving the phenomenal field not only interact with cortical nodes so as to create particular conscious states, they also function to produce a unified total state of consciousness. On this view, the 'enabling mechanisms' of creature consciousness are constitutively involved in the construction of consciousness. If this model is on the right track, then the science of consciousness ought to be as concerned with the mechanisms of creature consciousness as it is with uncovering the mechanisms underlying the contents of consciousness; that is, it ought to be concerned with both the neural states responsible for differentiating phenomenal states from each other, and the neural states that are common to each of the creature's phenomenal states. Nonetheless, state-based research paradigms are not fruitless, for by contrasting those conditions in which a certain neural node generates conscious content with those conditions in which it fails to generate conscious content we can explore the interface between the mechanisms that underlie the phenomenal field and those that underlie the contents of consciousness.¹³

References

Alkire, M.T., Haier, R.J., Fallon, J.H. 2000. Toward a unified theory of narcosis: brain imaging evidence for a thalamocortical switch as the neurophysiologic basis of anesthesia-induced unconsciousness, *Consciousness and Cognition*, 9: 370-86.

¹³ This paper was presented at the 11th annual meeting of the Association for the Scientific Study of Consciousness (Las Vegas), and I am grateful to the audience for their comments. I am also very grateful to Jakob Hohwy and Nicholas Shea for their comments on a previous version of this paper, and to previous discussion on this topic with David Chalmers.

- Alkire, M.T. & Miller, J. 2005. General anesthesia and the neural correlates of consciousness, In S Laureys (ed.) *Progress in Brain Research Vol 150: The Boundaries of Consciousness*, 445-55.
- Allison, T., Puce, A., Spencer, D.D., McCarthy, G. 1999. Electrophysiological studies of human face perception. I. Potential generated in occipitotemporal cortex by face and non-face stimuli. *Cerebral Cortex* 5: 415-30.
- Angel, A. 1991. The G.L. Brown lecture; Adventures in anaesthesia, *Experimental Physiology*, 76: 1-38.
- Angel, A. 1993. Cortical neuronal pathways and the process of anaesthesia, *British Journal of Anaesthetics*, 71: 148-63.
- Baars, B. 1998. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press
- Bayne, T. & Chalmers, D. 2003. What is the unity of consciousness? In A. Cleeremans (ed.) *The Unity of Consciousness: Binding, Integration, and Dissociation*. Oxford: OUP.
- Bartels, A. & Zeki, S. 1998. The theory of multi-stage integration in the visual brain, *Proceedings of the Royal Society of London Series B*, 265: 2327-32.
- Bayne, T. 2005. Divided brains and unified phenomenology: An Essay on Michael Tye's "Consciousness and Persons" *Philosophical Psychology*, 18/4: 495-512.
- Block, N. 2005. Two neural correlates of consciousness. *Trends in Cognitive Sciences*, 9/2: 46-52.
- Block, N. In press. The methodological puzzle of consciousness and the mesh between psychology and neuroscience, *Behavioural and Brain Sciences*, xxxx
- Carruthers, P. 2005. Conscious experience versus conscious thought. In *Consciousness: Essays from a Higher-Order Perspective*. Oxford: OUP.
- Chalmers, D. 2000. What is a neural correlate of consciousness? In T. Metzinger (ed.) *Neural Correlates of Consciousness*. Cambridge, MA: MIT Press.
- Crick, F. & Koch, C. 1990. Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, 2: 263-75.
- Crick, F. & Koch, C. 2003. A framework for consciousness, *Nature Neuroscience*, 6/2: 119–126.
- Dainton, B. 2000. *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.

- Dehaene, S. and Naccache, L. 2001. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework, *Cognition* 79: 1-37.
- Dennett, D. 1991. *Consciousness Explained*. Brown and Little.
- Driver, J. & Mattingley, J.B. 1998. Parietal neglect and visual awareness. *Nature Neuroscience* 1: 17-22.
- Epstein, R., Harris, A., Stanley, D. & Kanwisher, N. 1999. The parahippocampal place area: Recognition, navigation, or encoding? *Neuron*, 23: 115-25.
- Goldman, A. 1993. The psychology of folk psychology. *Behavioral and Brain Sciences*, 16: 15-28.
- Graham, D.I., Maxwell, W.L., Adams, J. Hume & Jennett, B. 2005. Novel aspects of the neuropathology of the vegetative state after blunt head injury. In. S Laureys (ed.) *Progress in Brain Research Vol 150: The Boundaries of Consciousness*, 445-55.
- Horgan, T. & Tienson, J. 2002. The phenomenology of intentionality and the intentionality of phenomenology. In D. Chalmers (ed.) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: OUP.
- Ishai, A., Ungerleider, L.G., Martin, A., Schouten, J.L., Haxby, J.V. 1999. Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences USA*, 96: 9379-384.
- Kanwisher, N. 2001. Neural events and perceptual awareness, *Cognition* 79: 89-113.
- Kanwisher, N., McDermott, J., Chun, M.M. 1997. The fusiform face area: a modular in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17: 4302-4311.
- Koch, C. 2004. *The Quest for Consciousness*. Englewood, CO: Roberts & Co.
- Laureys, S., Faymonville, M.E., Luxen, A., Lamy, M., Franck, G., & Maquet, P. 2000. Restoration of thalamocortical connectivity after recovery from persistent vegetative state, *Lancet* 355: 1790-1791.
- Laureys, S. 2005. The neural correlate of (un)awareness: lessons from the vegetative state, *Trends in Cognitive Sciences*, 9/12: 556-59.
- Laureys S, Owen AM, and Schiff ND, 2004. Brain function in coma, vegetative state, and related disorders. *Lancet Neurol*, 3: 537-46.
- Levy, J. 1977. Manifestations and implications of shifting hemi-inattention in commissurotomy patients. *Advances in Neurology* 18: 83-92.

- Levy, J. 1990. Regulation and generation of perception in the asymmetric brain. In C. Trevarthen (ed.) *Brain Circuits and Functions of the Mind: Essays in Honour of Roger W. Sperry*. Cambridge: Cambridge University Press, 231-48.
- Levy, J. & Trevarthen, C. 1976. Metacontrol of hemispheric function in human split-brain patients. *Journal of Experimental Psychology: Human Perception and Performance* 2: 299-312.
- Lewis, D. S. 1994. Reduction of mind. In S. Guttenplan (ed) *A Companion to the Philosophy of Mind*. Oxford & Cambridge MA: Blackwell (412-31).
- Logothetis, N.K. 1998. Single units and conscious vision. *Philosophical Transactions of the Royal Society of London Series-B Biological Sciences* 353: 1801-1818.
- Marks, C. 1981. *Commissurotomy, Consciousness and Unity of Mind*. Cambridge, MA: MIT Press.
- McCarthy, G., Puce, A., Gore, J., Allison, T. 1997. Face-specific processing in the human fusiform gyrus, *Journal of Cognitive Neuroscience*, 9: 605-10.
- Merker, B. 2007. Consciousness without a cerebral cortex: A challenge for neuroscience and medicine, *Behavioral and Brain Sciences*, 30/1: 63-81.
- O'Brien, G. and Opie, J. 1998. The disunity of consciousness, *Australasian Journal of Philosophy*, 76/3: 378-95.
- Passingham R.E. & Lau H.C. 2006. Free Choice and the Human Brain. In S. Pockett, W. Banks and S. Gallagher (eds.) *Does Consciousness Cause Behavior? An Investigation of the Nature of Volition*. Cambridge, MA: MIT Press.
- Pitt, D. 2004. The phenomenology of cognition Or What is it like to think that P? *Philosophy and Phenomenological Research*, 69: 1-36.
- Quiroga R.Q., Reddy R., Kreiman G., Koch C., Fried I., 2005. Invariant visual representation by single neurons in the human brain, *Nature* 435, 1102-110.
- Rees, G. 2001. Neuroimaging of visual awareness in patients and normal subjects, *Current Opinion in Neurobiology*, 11/2: 150-56.
- Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., & Driver, J. 2002. Neural correlates of conscious and unconscious vision in parietal extinction. *Neurocase* 8/5: 387-393.
- Rosenthal, D. 2005. *Consciousness and Mind*. Oxford: Oxford University Press.

- Schiff, N.D. 2004. The neurology of impaired consciousness: challenges for cognitive neuroscience. In: MS Gazzaniga (ed.) *The Cognitive Neurosciences* (3rd ed). Cambridge, MA: MIT Press.
- Schiff, N.D. & Plum, F. 1999. Cortical processing in the vegetative state, *Trends in Cognitive Science* 3/2: 43-44.
- Schiff, N.D. & Plum, F. 2000. The role of arousal and gating mechanisms in the neurology of impaired consciousness, *Journal of Clinical Neurophysiology* 17: 438-452.
- Searle, J. 2000. Consciousness, *Annual Review of Neuroscience*, 23: 557-78.
- Sheinberg, D.L. & Logothetis, N.K. 1997. The role of temporal cortical areas in perceptual organization. *Proceedings of the National Academy of Sciences (USA)* 94: 3408-3413.
- Stalnaker, R. 1984. *Inquiry*. Cambridge MA: MIT Press.
- Strawson, G. 1994. *Mental Reality*. Cambridge, MA: MIT Press.
- Tong, F., Nakayama, K., Vaughan J.T, Kanwisher N. 1998. Binocular rivalry and visual awareness in human extrastriate cortex, *Neuron* 21: 753-9.
- Tononi, G. 2004. An information integration theory of consciousness, *BMC Neuroscience* 5: 42.
- Tononi, G. 2007. The information integration theory of consciousness. In M. Velmans & S. Schnieder (eds). *The Blackwell Companion to Consciousness*. Oxford: Blackwell.
- Tononi, G. & Edelman, G.M. 1998. Consciousness and complexity. *Science*, 282: 1846-51.
- Tye, M. 2003. *Consciousness and Persons*. Cambridge, MA: MIT Press.
- White, N.S. & Alkire, M.T. 2003. Impaired thalamocortical connectivity in humans during general-anesthetic-induced unconsciousness. *Neuroimage* 19: 402-11.
- Zaidel, D. 1994. A view of the world from a split-brain perspective. In E.M.R. Critchley (ed.) *The Neurological Boundaries of Reality*, London. Farrand Press.
- Zaidel, E., Iacoboni, M., Zaidel, D.W., Bogen, J.E. 2003. The callosal syndromes. In K.H. Heilman and E. Valenstein (eds.), *Clinical Neuropsychology*. Oxford: Oxford University Press.
- Zeki, S. 2007. A theory of micro-consciousness. In M. Velmans & S. Schnieder (eds). *The Blackwell Companion to Consciousness*. Oxford: Blackwell.