Andrew Brook. *Atomism, Anti-representationalism and a Sketch of an Alternative.* Carleton University Cognitive Science Technical Report 2001-02 (www.carleton.ca/iis/TechReports). November, 2001.

Atomism, Anti-representationalism and a Sketch of an Alternative

Andrew Brook (abrook@ccs.carleton.ca), Carleton University.

This is the preface and first chapter of a book in progress on consciousness. If you would like to cite or quote from this paper, please check with the author first.

A Representational Theory of Consciousness

Andrew Brook

Table of Contents

Preface

Now is an interesting time to be working on consciousness. From the advent of cognitive science in the 1950s and 1960s up until the 1980s, few philosophers or other cognitive researchers did much with it. The philosopher Daniel Dennett once put the situation this way:

> Consciousness appears to be the last bastion of occult properties, epiphenomena, immeasurable subjective states – in short, the one area of mind best left to the philosophers. Let them make fools of themselves trying to corral the quicksilver of "phenomenology" into a respectable theory. [1978, p.149]

He could have added that this was pretty much the attitude of most philosophers, too, especially philosophers in the analytic tradition related to cognitive science. In the period in question, one could easily have concluded that for most cognitive researchers, cognitive functioning could proceed perfectly well without any consciousness at all.

This situation began to change in the mid-1980s. The work of the psychologist, Bernard Baars (especially 1988), was pivotal. He developed a methodology that he called contrastive analysis (compare the difference made by performing a task consciously and without consciousness). For the first time, researchers has a technique better than the widely discredited appeal to introspection to study consciousness and consciousness studies began to flourish. Centres of Consciousness Studies sprang up in a number of universities, the most famous being at the University of Arizona. A bit later, an email list called Psyche began and immediately attracted most of the significant researchers on consciousness as members. The Association for the Scientific Study of Consciousness was started: as of the summer of 2001, it had held its fifth international conference. And so on. Consciousness studies is suddenly playing a significant role in cognitive research.

In the years since 1988, two things have happened. There has been a huge explosion of work, with dozens of important books by philosophers alone. And there has been an explosion of new terminology and new theories – just what one would expect in a field where the subject is an intensely difficult one and serious multidisciplinary work is being done for the very first time. Here is some of the blossoming terminology: access consciousness, phenomenal consciousness, self-consciousness, simple consciousness, creature consciousness, state consciousness, monitoring consciousness, awareness taken to be coextensive with consciousness, awareness distinguished from consciousness, higher order thought, higher order experience, qualia, displaced perception, transparency .... and on and on. Except when I need to discuss the idea associated with one of these terms, I will mostly ignore then and develop terminology in my own merry way. The reason for doing so will be apparent very early in Chapter 1.

This book will strike most people as eccentric. It does not obsess about 'qualia' and other putative properties of individual psychological states, so it will strike most philosophers as eccentric (though I do say something about qualia). It does not obsess about attention as the basis of consciousness either, so it will strike most experimental psychologists as eccentric (though it eventually says a little bit about attention). Well, you can't please everyone! Neither qualia as usually understood nor attention captures what is distinctive about consciousness.

Thank yous. Acknowledgements.

Atomism, Anti-representationalism and a Sketch of an Alternative

1. Two Dominant Philosophical Views

Reading recent writings on consciousness by philosophers and by cognitive psychologists together can be a bit unnerving. Most philosophers focus on individual psychological states – individual perceptions or feelings or imaginings – or at most tiny combinations of such states (Rosenthal, 19xx; Chalmers, 1996; Dennett 1991 and P. M. Churchland 1995 are exceptions). Experimental psychologists by contrast focus on properties of whole cognitive systems. Attention is particularly favoured; for many of these researchers, to be conscious of something simply is to pay attention to it (Posner 1994, Mack and Rock 1998). Do all these people think that they are talking about the same thing?

Much of this book is devoted to working our way out of ideas dominant in recent philosophy to a position somewhat more like the position occupied by recent experimental approaches to consciousness (but only somewhat more like). When what philosophers are talking about is understood aright, the picture adopted by the experimentalists takes on a lot of plausibility. What supports experimental investigations trumps the 'thought-experiments' and other techniques of the philosophers. Surprise! To be sure, the techniques of the philosophers play an important role in research but these techniques have been considerably misapplied in recent consciousness studies.

We will start by trying to delineate some views dominant in recent philosophical work on consciousness. The first is a kind of atomism. Recent philosophers tend to talk about individual conscious states one by one ('what is it like for something to look red?') or at most in tiny groups. In both cases, the cognitive system that has them is entirely ignored. (Theorists may add, '...look red *to me?'* but they do nothing with the addition.) Consider as a leading example the massive literature on *qualia,* the felt quality of psychological states, states that 'it is like something to have', to use Thomas Nagel's (1974) now-famous phrase. The states purported to have qualia, what these qualia are like, the relationship of the qualia of states to representational content, and so on and so forth, are discussed without any mention whatsoever of thing that has the states. Let us call this view *atomism* about consciousness.

> **Atomism** – the view that conscious states can be studied one by one or in small groups, in complete isolation from the cognitive system that has them.

Such atomism is remarkable. It is obvious that consciousness does not come in atomically separable states in this way. Consciousness does not exist apart from conscious beings. It is not as though you could take individual conscious states away from the thing that has them and, for example, sell them! Turn to the experimental literature and one immediately find a better picture. The models developed by theorists such a Baars (19xx), Jackendoff (19xx), and Posner (1994) model conscious systems, not individual conscious states. Atomism has been so prevalent in the work of philosophers, however, that it has become a serious problem.

Atomism about consciousness goes with another view, which I will call *local realism*. Local realism is the view that consciousness or what is distinctive about consciousness, for example that in virtue of which it is like something to have a psychological state, is a local property of individual psychological states or tiny groups of psychological states. By 'local property' I a mean property that is either a nonrelational property of single psychological states or, though relational, ties only very small groups of psychological states to one another. A relationship between one state and another single state would be an example. This account of a local property is not very precise but it is good enough for my purposes. It is the contrast with nonlocal properties that matters; by 'nonlocal property' I have in mind the kind of properties that are those central to theories that view consciousness as a relationship between a great many psychological states and a conscious being whose states they are. And *local realism* is the view that consciousness or what is distinctive about consciousness is either a non-relational property of individual psychological states or a relationship among very small numbers of psychological states.

> **Local realism** – the view that consciousness or what is distinctive about consciousness is either a non-relational property of individual psychological states or a relationship among very small numbers of psychological states.

The view is local because it ascribes consciousness to individual states or tiny groups of states in isolation. It is realist because it takes these states to be real (not illusions that we believe in only because of certain concepts that we have, or products of interpretation, or something like that).

There are (at least) two types of local realism. In one type, appearing red to me would be a property of an experience of red, being painful is a property of a pain I am having, and so on. In the other type, a representation of red gets to be conscious by being related to another psychological state, for example by becoming the object of a thought about that representation (the so-called higher-order thought view of consciousness [Rosenthal 19xx].)

It might seem that atomism requires local realism but in fact that is not clear. Some atomists about consciousness simply say nothing about whether qualia, for example, are local or nonlocal properties of the states they discuss. This neutrality is a bit curious because these theorists believe that they can say other important things about qualia, e.g., that when it is like something to have a representation, this quale, this being like something, is radically different from other aspects of the representation, but neutral they have been. Nonetheless, local realism would certainly promote atomism: if consciousness is a local property of certain states, it would be very tempting to hold that one could study such states one by one and in isolation from the system that has them.

As we have already said, at least the local part of local realism is almost certainly not true. The way it is not true puts a lot of pressure on atomism. Consciousness is not a (local) property of individual psychological states. Certainly we are *conscious of* various things: the world around us, our bodily states, our psychological states (including the representational states that make us conscious of these other things), and ourselves. But are any of the things that we are conscious *of* themselves conscious states (in any way other than being states that we are conscious of)? That is far from clear. It is like something *for me* to have these states. These states and what they represent appear *to me* as this or that. But is this being like something to have, this appearing in a certain way, a property of the state that thus appears? At most consciousness seems to be a complex relational

property connecting a state to me, the being who is conscious of it. This observation by itself is enough to create some problems for contemporary philosophical work on consciousness.

It is important to note that undermining *local* realism about consciousness is not the same as undermining *realism* about consciousness. Even if consciousness is not a local property of individual psychological states, it could still be a genuine property of cognitive systems, a property that is instantiated in a form recognizable as consciousness. I mention this now because there have been influential treatments of consciousness recently that back off from any form of cognitive-system realism about consciousness, for example Davidson's (19xx) view. In Davidson's view, consciousness arises out of a complex triangular interaction among oneself, other purposive beings, and the world. By itself, this triangulation picture need not depart from realism; the result of the triangulation, consciousness, could still be a real property of cognitive systems. For Davidson, however, not only does consciousness arise out of triangulation, it is (roughly) nothing more than triangulation. When triangulation results in stable attributions of consciousness to self and others, that is what consciousness is. And *this* view *is* incompatible with most versions of realism about consciousness.[1]

Denying that the property of being conscious is a local property of individual psychological states and insisting that it is a property of cognitive systems as a whole immediately invites the question: What, then, *is* consciousness? A good question and one that will not receive a final answer in this book for quite some time. At minimum, being conscious of something is one way of representing something.[2] Of course, things can also be represented unconsciously. In fact, probably the vast bulk of our representations never make us conscious of anything, not in the world and not

---

[1.] In these introductory remarks, I want to keep the discussion relatively uncluttered. For those who know the literature and would like to see how my discussion relates to other work, I will include some pointers in footnotes. For those just interested in my view, the footnotes in this chapter can be safely ignored. The connection I want to point to here goes like this.

Some have thought that Dennett's (1991, 1996, 1998) treatment of consciousness rejects realism about consciousness, too. Dennett vehemently denies this and I think he is in most respects right. However the issue is complicated. We will touch on it again later in this chapter and consider it in detail in Chapter 2.

[2.] The relationship between being conscious and representing may be more complicated than at first appears. Often when we are conscious of something, we are representing that thing. For example, when we are conscious of something in the world around us or, in proprioception, a state of our own body, we are representing that thing. But what about the case where what we are conscious of is the experience, the representation, itself, rather than what it is representing? There is a disagreement in the literature about what is being represented here. The majority view is that when one is conscious of one's representations, one is representing the representation, i.e, something different from what representation represents. However, some theorists argue that consciousness of one's representations is also given merely by representing something in them. Focus on a square red patch. Now focus on your *awareness* of that square red patch. Are you now aware of new properties? Philosophers who answer 'no' include Harman (1990), Dretske (1995), Tye (1995), and Crane 19xx). The resulting theories are sometimes called displaced perception or transparency theories. For now we do not need to go into this issue. However, it will loom large in Chapter yy.

in our own minds. A representation does not need to make us conscious of something to be cognitively active.

One major issue about the nature of consciousness is whether the difference between conscious and nonconscious representation can be captured by appealing to representational properties of conscious states alone or whether the properties that make a state conscious are nonrepresentational properties. Here atomist, local realist researchers split into two camps. Representationalists about consciousness maintain that consciousness is a representational properties of conscious states. Anti-representationalists maintain that it is not. For them, the difference between a state that is conscious and one that is not is *not* a difference in how that state represents anything or a difference in the kind of representation it is or anything else representational. Here is how the anti-representationalist view arises.

When something appears to us to be a certain way, the representation in which it appears that way can play two roles in our cognitive economy. On the one hand, the representation (or the contents of the representation) can connect inferentially to other representations: if the stick appears to have two straight parts with a bend in the middle, this will preclude representing it as forming a circle. The representation can also connect to belief: if the stick appears straight with a bend in it, I will not form a belief that it bends in a circle. And to memory: I can compare this stick as it appears to sticks I recall from the past. And action: if I want something to poke into a hole, I might reach for the stick. In all these case, so long as I am *representing* the stick in the appropriate way, it would seem to be irrelevant whether I am *conscious* of the stick or not. My representation could do these jobs for me just as well even if I were not aware either of the stick or of my representation of it. But I am also *aware* of the stick – it does *appear* to me in a certain way. This seems to be something different from any representational properties of the representation, at any rate properties such as those we just considered.[3]

[3.] Chalmers' well-known (1995) distinction between what he calls the easy problem and the hard problem of consciousness starts from this distinction between the cognitive role of representations and something appearing to be like something in them. Understanding the two phenomena would appear to be quite distinctive problems. Chalmers calls them the easy problem and the hard problem. The easy problem is to understand the inferential and other roles of such states. The hard problem is to understand how, in these states or any states, something could appear as something to me, how certain stimulations of the retina, processing of signals by the visual cortex, application of categories and other referential and discriminatory apparatus elsewhere in the brain can result in an *appearing,* a state in which something *appears* a certain way. Chalmers says that the easy problem is easy because it is simply the problem of the nature and function of representation in general, while the hard problem is hard because it appears to be *sui generis,* quite unlike any other problem about cognition that we face. If the first problem is easy, I'd hate to see what a hard one is like but there do seem to be two distinct issues here and the issue of how anything can appear to us at all does seem to be special.

One element in what makes the hard problem appear special is Levine's (1983) explanatory gap. According to Levine, not only do we not know of any mechanism or causal process that would allow us to understand the phenomenon of things appearing to us in terms of information processing or brain processes, we cannot even imagine what such a mechanism might be like. There is nothing like the same explanatory gap with respect to cognitive functioning.

Since there are atomists and local realists about consciousness who are anti-representationalists (Block 19xx, 19xx, Chalmers 1996, 19xx) and others who are not (Lycan 19xx, 19xx, and Tye 19xx, 19xx), it is worth flagging anti-representationalism explicitly:

> **Anti-representationalism** – the view that the difference between a state that is conscious and one that is not is not a difference in how that state represents anything or a difference in the kind of representation it is or anything else representational.

We will examine some arguments for the view in Chapter 2 and others in Chapter 7. These arguments often take the following form: the felt quality of a state could change while its representational properties remain the same. The arguments are usually based on thought experiments. One well-known argument is the inverted spectrum argument: how colours appear to us could be inverted without changing how our representations of colour function as representations. Another is the zombie argument: there could be creatures for whom it is not like anything to represent anything and yet their representations function cognitively as representations function in us.

Sometimes such arguments go as far as to conclude that what is distinctive to consciousness is not only not representational, it is not even physical. One way of arguing this to make one's zombie a microphysical duplicate of oneself. If a zombie such as this is possible, then qualia are not a physical property of me. Another famous argument to the same conclusion is Jackson's (19xx) Mary, the colour scientist who knows everything there is to know about the experience of colour, therefore everything physical there is to know about the experience of colour, but who has never experienced colour herself. Then she experiences colour. Clearly she gains something she did not have before. However, she knew everything physical about colour. Therefore, what she gains must be something nonphysical.

It is not clear that any of these thought experiments establish real possibilities, or, if they do, entail the conclusion drawn from them. We will examine them in Chapter yy. The point for now is this. Because of Occam's Razor, the representational theory of consciousness stands or falls on the success of the anti-representational case. All parties agree that representations exist, so if one party wants to argue that there is in addition something nonrepresentational, the onus is on her to make the case for this additional element. If there is no reason to believe that such an additional element exists, we have an excellent reason to believe that it does not. Anti-representationalists maintain that there is something about states of which we are conscious, the felt quality or quale of them, that is different from anything that makes a representation a representation. Advocates of the representational theory urge that the opponents of representationalism are wrong about this. For good or ill, anti-representationalism is built on arguments like the ones just sketched. If they don't work, anti-representationalism is in deep trouble. Since I favour a version of the theory that consciousness is a representational property of certain psychological states, that is just what I will be arguing.

So the book has three major targets: atomism, local realism, and anti-representationalism about consciousness. What does the alternative point of view look like? The view I advocate is that consciousness is:

- a property of whole cognitive systems,
- a real property of such systems, and,

- a representational property of such systems.

We will sketch such a model in Section 5. Before we can do that, we have to do some more spade work.

2. The Diversity of Consciousness: Is there a Common Core?

In the recent literature on consciousness, there are disagreements even more basic than the ones just sketched over atomism, local realism, and anti-representationalism. Return to our question, 'Do all these people think that they are talking about the same thing?' In fact, it is not clear that theorists *are* always talking about the same thing. There is disagreement even about something as basic as what the term 'consciousness' refers to. Many theorists take it to be about a kind of access to one's psychological states, where the contrast is perhaps with unconscious psychological life of the sort that Freud talked about. On this construal, most animals are not conscious. Others take 'consciousness' to be broad enough to include access (the right kind of access) to one's world, where the contrast would be the nonconscious, not the unconscious. On this proposal, we then need a further term for the special access each person has to herself and her psychological states: consciousness of self or self-consciousness. The duality is brought out nicely by a statement such as, 'Animals have conscious states but are not conscious of having them'. Some theorists treat both such states and consciousness of having them as kinds of consciousness. Some theorists restrict the term 'consciousness' to consciousness of one's states. In fact, the latter is the most common use of 'consciousness' philosophers and cognitive psychologists. On this usage, the term 'consciousness' refers to consciousness of oneself and one's states ("the subjective character of our mental states," as Perry recently put it [2001, 45]), though researchers are often not very clear about how they are using the term. Here is how I see the matter.

We use the term 'conscious' (or 'aware'; in most usages, 'aware' is simply an alternative term for 'conscious') and cognates to speak of a wide variety of things:

- Perception and proprioception. 'At that moment, the dog became conscious of the cat.' 'I just noticed that I'm crossing my legs again. Blast!' (said by someone who has just has an operation to repair a torn knee ligament).
- One's psychological states. To say 'I am feeling less anxious than I was yesterday', you must be conscious of the anxiety. (Here I have deliberately chosen an example where there is no representation of anything, not overtly at any rate.)
- Oneself. To say 'I am not pleased with myself for forgetting that appointment', you must be conscious of yourself. That's the object of your displeasure.
- Something that can come in degrees. 'I was vaguely conscious of something coming up on my left, so I turned to get a good look at it.'
- A global state of certain beings. 'She is slowly regaining consciousness.'

One response to this diversity would be to say that we simply use the word 'conscious' and cognates to refer to a bunch of different things that have little if anything in common. To resolve the

confusion, we should confine the word to one of these kinds of things and use other terms for the others, or even drop the term altogether, replacing it with a number of more precise terms.[4] The trouble with this proposal is that, intuitively, the things being talked about in the five examples do seem to have something in common. They do not seem to be just a bunch of unrelated phenomena. Is there a common core to our different uses of the word 'conscious' and cognates?

3. A Common Core

Part of what is common to our five examples is that in every one of them, we are talking about how something *appears* to us, seems to us: how things in the world appear to us, how our own states appear to us, how one appears to oneself.[5] The only exception is the last example, regaining consciousness. It is only a partial exception because here we talking about regaining our *capacity* for things to appear to us. This is the first part of a common core:

> **Common core of the concept of consciousness, first move:** consciousness is things appearing to us and various capacities for things appearing to us.

This point also gives us part of what we need to clear up the muddle in the literature about what 'consciousness' refers to. Whether it is consciousness of the world or consciousness of oneself and one's psychological states, when we talk about consciousness, we are talking about something appearing to us in some way. How something appears to us either is or is a property of how we are representing the thing, so whether one has in mind consciousness of the world or consciousness of oneself and one's states, one is talking about a property of representings of things, not of states of affairs represented. If so, whether a theorist is talking about consciousness of the world, consciousness of oneself and one's states, or consciousness as a global property of a cognitive system, one is talking about the same *kind* of thing: how something appears as one is representing it, and the abilities that go with this.

> **Common core of the concept of consciousness, second move:** consciousness is things appearing to us *as we are representing them*, and various capacities for things appearing to us.

The next stage in clearing up the muddle is to bring into focus what philosophers are in fact talking about when they restrict consciousness talk to consciousness of one's own psychological states (whether or not they know that this is what they are doing). Start with what is left out here, namely, consciousness of the world. (Consciousness of one's own bodily states such as limb position is often left out, too.) When we talk about simply being conscious of the world, we are talking about the world appearing in various ways to us. In this kind of consciousness, the thing that is conscious

---

[4.] This is the kind of proposal that one might expect an eliminative materialist to make, for example Patricia Churchland (1983xx).

[5.] The Kantian allusion in the term 'appearance' is no accident. As will become clear, the theory I advocate is broadly speaking a Kantian one (Brook 1994).

of the world need not be *conscious* of how the world appears to it. Think of mammals such as dogs and cats. These creatures, it is plausible to hold, are conscious of the world (and of states of their body). These things appear to them in various ways. There is little reason in most cases to think, however, that they are conscious of *how* these things appear to them or even *that* they are appearing, i.e., how or even the fact that these things appear to them does not themselves appear to them. Put another way, they are conscious of the world (and states of their own body) but not of their experiences of these things.

The distinction I am trying to draw between consciousness of the world, i.e., things (in fact) appearing to one, and consciousness of how the world appears to one can be illustrated by some recent work in cognitive neuroscience.

Blindsight is often invoked in contexts like the current one. Due to damage to the visual cortex, blindsight patients have a scotoma, a 'blind spot', in part of their visual field. Ask them what they are seeing there and they will say, "Nothing". However, if you ask them instead to *guess* what is there, they guess with far better than chance accuracy. If you ask them to reach out to touch whatever might be there, they reach out with their hands turned in the right way and fingers and thumb at exactly the right distance apart to grasp anything that happens to be there. And so on (Weiskrantz, 1986). I am not going to invoke blindsight, however. I am going to use what is for our purposes industrial strength blindsight, what is tendentiously called 'inattentional blindness'. (Tendentiously because there is in fact a huge debate raging about whether the phenomenon in question has anything to do with attention [Mack, http://psyche.cs.monash.edu.au/v7/psyche-7-16-mack.htm].)

In inattentional blindness research, the subject fixations (concentrates on) on a point and is asks to note some feature of an object introduces on or within a few degrees of fixation.

[Figure 1]

After a few trials, a second object is introduced, in the same region but clearly distinct from the first object.

[Figure 2]

Subjects have no inkling that a second object will appear. When the appearance of the two objects is followed by 1.5 seconds of masking, at least one-quarter of the subjects or, if the parameters are set right, almost all subjects have no awareness of having seen the second object. Yet – and this is what makes inattentional blindness better for our purposes than blindsight – when the second object is a word, subjects clearly encode it and process its meaning. Evidence? When asked shortly after to do a 'stem completion task' (i.e., to complete a word of which they are given the first two or three letters), they complete the word in line with the word they claim not to have seen much more frequently than controls. In inattentional blindness, subjects are just as blind, consciously, to the information before them as blindsight patients are. However, their informational access to the word they don't see is very complete.

This distinction between an object appearing to a subject, as evidenced by the subject processing up to semantic information provided by it, and yet the subject having no consciousness of taking in the information is the distinction I am trying to make. Whereas with blindsight patients barely able to do anything with the information they access 'in' the scotoma, some resist saying that

they have any consciousness of the items on which they cannot report, with inattentional blindness it is natural to say that they do, that the inattentional blind are about as well of as most animals without the ability to self-report. The kind and degree of sophistication of the discriminations and other uses of the incoming information that these subjects make is similar. Yet it is not like anything for these subjects to have this information.[6]

Recent discoveries about the ventral and dorsal streams in the brain (Milner and Goodale, 1995) present another nice example of the distinction. The evidence Milner and Goodale have amassed strongly suggests that the ventral stream supports the processing involved in consciousness of the world, specifically visual consciousness, while the dorsal stream supports control of fine motor action. It turns out that there is a double dissociation between the two, as patients with damage to either stream but not both streams demonstrate. Patients with certain kinds of damage to the ventral stream lose visual consciousness of what is before them but continue to be able to control fine bodily movement with respect to the items before them. Other patients with certain kinds of damage to the dorsal stream lose to the ability to control fine bodily movement but retain visual consciousness of what is before them.

In fact, the dissociation can be displayed in normal subjects. Consider the Titchener illusion, in which a circle appears bigger or smaller depending on whether it is surrounded by smaller or bigger circles.

[insert illustration of Titchener illusion]

When the illusion is set up using circular pieces of plastic rather than drawings on a sheet of paper, subject for whom the circles *appear* to be different sizes will still, when they are asked to reach for the circles, open their fingers to exactly the same size, and the right size, for both circles (Milner and Goodale, 1995, p. 168)! A remarkable finding indeed!

The interpretation of Milner and Goodale's findings is controversial but one natural way to interpret them is to say that ventrally-damaged patients retain at least enough consciousness of their

---

[6.] Change blindness, attentional blink, and visual neglect are related phenomena. In change blindness, subjects fail to notice some obvious change. Sometimes the change is made during a saccades of the eyes (saccades are rapid movements of the eyes; while the eye is moving, information on the retina does not 'register'). Sometimes the change is merely unexpected. For example, a researcher unknown to a student engages the student in conversation. Two other students carry a sheet of plywood between them. The researcher switches places with one of the people carrying the plywood. The student almost never notices that he is talking to a new person! The attentional blink is

Visual neglect is a strange result of certain kinds of brain damage in which people cease to be aware of, for example, one side of their visual field or one side of their body.

What ties all these phenomena together for our purposes is that in all of them, the information not being registered consciously nevertheless is encodes and enters into cognitive tasks making use of semantic information, for example, disambiguation of ambiguous sentences.

world for motor control purposes but lose consciousness of how the world appears to them and vice-versa for the dorsally-damaged. The latter retain consciousness of how the world appears to them but lose the ability to use the information thus accessed for motor control.

As much as possible, I am avoiding discussion of the work of philosophers in this introductory chapter or relegating such discussion to footnotes. However, a distinction of Ned Block's has become so well-known and so influential and might appear to be so directly relevant to the distinction that I am making that I have to say a word about it. I have in mind Block's (1995) distinction between P-consciousness and A-consciousness. 'P' stands for 'phenomenal' and 'A' stands for 'access'. Block defines 'A[access]-consciousness' like this: "A state is A-conscious if it is poised for direct control of thought and action" (1995, p. 233). He tells us that he cannot define 'P-consciousness' in any "remotely noncircular way" but will use it to refer to what he calls "experiential properties", what it is like to have certain states (1995, p. 231). What I am calling consciousness of the world may appear to be rather like Block's A-consciousness. It is not. It is a form of P-consciousness. Consciousness of an object is – at this stage we can say only – *consciousness* of the object.

Indications that what we are calling consciousness of the world is some form of *consciousness* include the following. Often the organism refers or points to the items in question. It focuses on these objects. The objects can increase the organism's level of alertness, especially the level of alertness aimed at the objects themselves. Often ensuing behaviour is behaviour appropriate, not to the way the objects actually are, but to how the objects looked (or could have looked) to the organism (Dennett, 1978x). Even if consciousness turns out *just to be* informational access of a certain kind (something that Block would deny), there would still be informational access without consciousness and informational access that is consciousness. In my notion of consciousness of the world, I am talking about the latter. Yet what I am talking about is clearly not P-consciousness either – beings who have it are not conscious of how things appear to them, i.e., when you are conscious of something in this way, it is not like anything to be conscious of it. Between simple informational access and P-consciousness, there is a third thing: consciousness of the world.

What is the difference between conscious access to the world and nonconscious informational access to the world? Think of what we know of such consciousness and how we know it. Our evidence that something is conscious of the world around it is almost entirely behavioural. Think of a dog regaining consciousness after a nap. Its regaining consciousness is a matter of it again making rich, supple discriminations and using the information thus gained for complex behavioural control, especially behaviour that relates to its relationships to us. Behaviour that requires us to invoke how the world is appearing to the animal is particularly relevant. Even when the animal appears to be expressing psychological states rather directly, as in crying out in pain or desperately trying to find something to drink when thirsty, our evidence for these states is entirely behavioural.

There is, of course, a great deal more to be said about animal consciousness than this (see, for example, ... 19xx) but what I have said is enough to show how consciousness of the world is different from Block's A-consciousness. There are all sorts of situations in which information exercises control over thought and action, yet no one would dream of calling the access conscious. Think of the effects of the information in ethanol on thought and action. (In a couple of good senses of 'information', chemicals contain information. They have structure and this structure can be used

to reduce uncertainty.) This is informational access but no one would call the access that the brain has to the information in ethanol a form of *conscious* access.

Enough about the distinction between consciousness of the world and consciousness of how the world (and other things such as experiences) appears to one. Let us draw out implications conclusions for the topic with which we began the section, namely, what philosophers are talking about when they use the word 'consciousness'. As we can now see, generally speaking they are talking, not about things appearing to something, but about *consciousness* of how things appear to one. Even if they ignore consciousness of the world, i.e., the world appearing to one without one being conscious of how it appears, and generally they do, they usually don't ignore *consciousness* of how the world appears. In sum, philosophers generally restrict their use of the word 'consciousness' to consciousness of how things appear to one.

Problems can arise when one does not notice that, in addition to consciousness of how things appear to one, we talk and talk perfectly sensibly about consciousness of the world. Consider this passage from Dennett:

> In order to be conscious – in order to be the sort of thing it like something to be – it is necessary have a certain sort of informational organization ... [one] that is swiftly achieved in one species, ours, and in no other. ... My claim is not that other species lack our kind of *self*-consciousness. ... I am claiming that what must be added to mere responsivity, mere discrimination, to count as consciousness *at all* is an organization that is not ubiquitous among sentient organisms [Dennett 1998, p. 337].

There is a lot to like in this quotation, in particular the idea that special, fairly sophisticated cognitive abilities are involved in achieving the kind of consciousness in which it is like something to have one's states, be oneself. We will return to this idea. Dennett gets his conclusion that few if any species other than human beings have consciousness, however, only by *restricting* the term 'consciousness' to consciousness of how things appear and *ignoring* the perfectly sensible usage in which we talk about consciousness of the world. Dennett might be right about the 'something it is like to have' kind of consciousness. That kind of consciousness might make cognitive demands that only human beings can meet. But there is another kind of consciousness still available to non-language-using animals.

If consciousness is a matter of things appearing – that they appear and consciousness of how they appear – then consciousness is a property (not necessarily a local property) of representings, not of states of affairs represented. Many complex questions have been asked about the nature of acts of representing. Here I will take up just one of them, because it has been thought to contain another argument against representational theories of consciousness. The issue is externalism about representational content. Many philosophers now maintain, to paraphrase Hilary Putnam (1975), that the content of representations ain't in the head. Philosophers who accept this view then go on to say one or the other of two ways about consciousness. Some continue to hold the commonsense view that qualia, the felt quality of representations, are in the head and, since representational content is not in the head, conclude that qualia are not representational content. This is the additional argument against representationalism I mentioned. Others hold that if representational content ain't in the heard, then how something appears ain't goin' to be in the head either (for this way of reacting to

externalism about content, see Tye, plato.stanford.edu/entries/qualia, p. 10). We will explore the implications of externalism for representational theories of consciousness in Chapter yy.

4. Easier and harder problems of consciousness

It is now common, thanks to David Chalmers (1995, 1996), to distinguish easy and hard problems of consciousness (see note 3). We can apply that distinction in a different way here. Of the varieties of consciousness that we have introduced, namely, consciousness of the world (and bodily states), consciousness of how things appear in one's own psychological states, and consciousness of oneself, consciousness of the world and of one's bodily states seems likely to be easier to understand than the other two. With consciousness of the world, a behaviourally-based information-processing approach is apt to take us a long way.

The going gets tougher when we move from consciousness of the world to consciousness of how the world *appears to us*. For example, the question, 'What it is for a shade of colour to appear to one in a certain way, for the experience of that shade of colour to be like something to have?' will probably be trickier to answer than the question, 'what it is to be conscious of a colour?' Part of the reason is that, in line with the suggestion in the passage from Dennett that we just considered, more is likely to be involved cognitively in being conscious of how something appears to us than in simply being conscious of the thing. In the former case, as well as the cognitive abilities involved in discriminating and characterizing things in the world, we would need the additional cognitive abilities involved in becoming conscious of properties of our own psychological states.[7] Clearly we can make progress with the question about consciousness of objects by talking about information flows, discriminatory powers, behavioural control. It is not clear what kind of information-processing talk would give us a handle on what it is to be aware of how things appear to us. Anti-representationalists about consciousness maintain, of course, that no information-processing talk is going to get to the heart of the matter here. Whatever the truth about that claim, clearly consciousness of how things appear to one is apt to be a harder problem than consciousness of the world.[8]

[7.] A fair amount of ink has been spilled over the question of what these additional cognitive skills might be like. See Peacocke 1993, Dretske 1995, 1997, Lycan 1997, Tye 1997. We will return to the question in Chapter yy.

[8.] What I am here calling the harder problem of consciousness is very similar to what Chalmers calls the hard problem (see note 4) but that is not true of what I am calling the easier problem. In our hard problems, we are both talking about the phenomenon of its being like something to be in certain states. Hence the allusion to his distinction in the title of this is not too misleading. However, what I am calling the easier problem is different from his easy problem. His easy problem has do with understanding the information flows in and about and the cognitive functions of conscious states. My easier problem is the problem of understanding what consciousness of the world and one's bodily states is like.

5. What is Appearing Like? Step 1: The Need for a Conscious Subject

As we saw, many recent analytic philosophers of mind have adopted an atomist, local realist approach to consciousness of how things appear. Quite a number of such philosophers are also anti-representationalists. What might a non-atomistic, realist but not local realist, representational model of things appearing to us be like? Here I will sketch one such model, one involving five steps. Much of the rest of the book will be devoted to filling out this sketch.

Actually, what I just said is not quite right. We will actually take the first step. The first step to a nonatomist, realist, representational alternative is to notice that consciousness is a matter of *something being conscious* of something. Consciousness requires a subject, a thing that is conscious. Return to our examples of consciousness talk. The second example introduces what philosophers are talking about when they talk about qualia: a property of certain psychological states. Certainly such states should enter a theory of consciousness. Notice, however, that what is being talked about in the last two examples is quite different. The properties in the last two examples are properties of a cognitive system as a whole, properties in which we are in that special relationship to various items that we call being conscious of them. Utterly unlike qualia as usually conceived, these properties rest on abilities of the system as a whole: an ability to be conscious of things, and an ability to be differentially conscious of things (vaguely conscious of this, clearly conscious of that).[9] These examples are about the cognitive system being conscious – of the world, of its own states, of itself – not anything to do with, or at any rate not anything confined to, individual psychological states.

It is plausible to think, moreover, that when certain states appear to me in some way and I am conscious of them, how they appear and what they, their contents, etc., appear to be like is the result of my exercising certain abilities, specifically, abilities to interpret, to judge. Something appearing to me in some way is not a nonrelational property of or anything 'supervening on' individual psychological states. And what am I interpreting or judging? It is plausible to say: information incoming in representations. If so, something *appearing* to me in some way has to be more than a nonrelational property of individual psychological states.

The subject is central to cognitive life in general, nonconscious as much as conscious. In general, we have no choice but to conceive of anything we describe in intentional language as involving a subject. (Intentional language is language for psychological states that are about something, in the way that a perception is about whatever is perceived, a belief is about whatever is believed. In philosopher's jargon, this property of aboutness is called *intentionality* and language that ascribes it is called *intentional language*.) If we characterize an event or state of affairs in the language of intentionality, we must postulate that it is had by someone. As Dennett has put it, "Wherever a theory relies on ... intentionality, there a little man is concealed" (1978a, p. 12). Every representation *of* something represents it *to someone.* Every intention is intended *by someone.* Every

---

[9.] A well-known response of David Lewis's (1990xx) to an even better-known thought experiment of Frank Jackson's about Mary the colour scientist is called the ability hypothesis. We will discuss this thought experiment in Chapter yy. For now I am making no claims about any relationship between my invocation of the notion of ability and Lewis's hypothesis.

action is the action of *some agent.* Everything meant by an utterance is meant *by someone.* Similarly, if we ascribe consciousness, we have no choice but to ascribe a subject of consciousness, a being that is conscious. There can be no awareness of something that is not *someone's*.

If states were conscious intrinsically, that is to say, would be conscious in the way that they are no matter where they were, how they were connected to other representations, or whether they were related to minds, then this would not be the case; consciousness would not require a subject of consciousness. The classical empiricist theories of representation of Locke and Hume may have maintained such a view. But there is a well-known, compelling reason to deny that any psychological state can be intrinsic in this sense; any psychological state can appear in different ways in different contexts and times, depending on how it is taken. A 'taking' requires a 'taker'. In general, anything can be conscious only in a system that is conscious. By accepting this requirement, psychological theories of consciousness are decisively superior to atomist philosophical ones.

As Dennett (1978, pp. 101, 122) and others have insisted, we must give an account of the conscious subject. It cannot be left as an undischarged homunculus or exempt agent (exempt from capture in a theory).[10] Here representational theories seem to face a problem. If consciousness of representations *presupposes* a subject, it is hard to see how they could be used to give an account of it, to 'discharge' it. As we conceive it, the conscious subject is something different from the states of which it is conscious, something able to take some states up, drop others, store yet others, without itself changing in any way that is significant to its being a conscious subject and continuing to be the subject it is. So, suggests Dennett hyperbolically, psychology is impossible! I am not so sure myself that representational theory is impotent on this matter; indeed, the main aim of this book is to present a representational theory of the conscious subject that can meet the requirement.

Instead of attempting to meet this requirement, even those analytic philosophers who have considered it seem to take the attitude that if they can just develop an adequate account of something else, the problem of the conscious subject will go away. Consider in this regard eliminativism and homuncular functionalism.

The eliminativists approach the problem of the subject by attempting to dispense with the intentional language that forces it upon us. (This is not to say that the need to postulate a subject does not arise in other ways, just that recourse to intentional language forces the issue.) Switch to the vocabulary of neuroscience to describe what we now call consciousness, as we should do for other reasons in any case according to them, and the need to postulate a conscious subject will disappear. By contrast, homuncular functionalists accept that we are stuck with intentional discourse. If we keep intentional discourse, we must also keep the subject, the little man, of course, and they are prepared to do so. However, they insist that we need not be stuck with one big, smart, unanalyzed homunculus. Rather, the big, smart homunculus can be reduced to a fleet of small, stupid homunculi flying in loose formation, each with a single or at most a small number of well-characterized functions. (Dennett, 1987, pp. 37-42, 1991, pp. 27-41; Lycan, 1990, Ch. 4 are good examples of

[10.] The need to discharge the homunculus was articulated as long ago as Leibniz; it is far from being a new issue.

homuncular functions.[11] Paul (1979) and Patricia (1986) Churchland are often cited as leading example of eliminativist. However, with respect to consciousness, it is not clear that they are. I will comment on this issue in more detail in Section 7.)

Unfortunately, the mind is more complicated than is dreamt of in either philosophy. It would be silly to expect everything in our initial picture of the conscious subject to survive in an completed account of cognition and consciousness but we do want an account of some of the things in that picture. And many of the things about which we want an account are properties of the whole conscious system, properties that strongly resist capture as some combination of stupid little homunculi. Some examples:

- We are aware of whole groups of representations at the same time.
- Often when we are aware of whole groups of representations, we are also aware of ourselves as the common subject of these representations.
- Judgments are not just made in us; sometimes we make judgments.
- We have a great many cognitive faculties that are available throughout cognition and some of them are closely linked to consciousness. Memory is a clear example. Language may be another.
- If various psychological functions are performed by different systems in the brain, there have to be processes that synthesize these various activities into the single, simultaneously-introspectible patterns of representation, belief and behaviour so central to our conscious life.
- Information being located in a conscious system is not enough; the information must be of use to it. What kind of access on the part of the conscious system does this require? (As we will see when we get to Chapter 8, this phenomenon of information being of use to a system is closely related to consciousness.)
- What is consciously controlled attention like and what is its role in accessing and integrating information?
- Consciousness is independent of, indeed can continue perfectly well in the absence of, sensory inputs
- Consciousness as a whole can disappear in deep sleep, and . . .
- reappear as whole though perhaps not entirely in dreams.[12]

[11.] Dennett's most heroic attempt to discharge the little man is *Consciousness Explained* (1991). He attempts to show how at least short-lived conscious subjects could result from competition among multiple, draft narrative-fragments. He calls them the `virtual captains'. This theory is a version of homuncular functional. We will examine it in Chapter yy.

[12.] This list started from but goes beyond Churchland's seven requirements on a theory of consciousness in (1995), pp. 213-14.

We demand an account of such phenomena. Some of them will be addressed in Chapter 3, others not until Chapter 8. The point here is that it is hard to see how they could be accounted for either by eliminativism or by homuncular functionalism.

If we recognize that the conscious subject cannot be ignored, we also cannot prejudge what it might be like. The someone to whom a representation is represented could in principle turn out to be like almost anything. It could even turn out to be a kind of representation.

6. The Other Four Steps

Step 2: The Role of the Conscious Subject in Shaping How Things Appear to It

Having taken step 1, we will now just sketch the other steps. Taking them is one of the jobs of the rest of the book. We want a model that, though realist, avoids local realism and models consciousness as representational. The next step is to show that when something appears to one to be like something, this is not anything that could be 'read off' a representation. It is the result of complex cognitive processing. When a representation or something represented in it to be like something to have, one has processed that representation and perhaps other things in various ways. We will introduce the ways we thus process and the extent of the processing in Chapter 2 and return to it in more detail in Chapter YY. For now, keep in mind that if something appearing to someone is a relationship between a conscious being and the state that appears, it could not be a local property of what appears.

Given this analysis, the question for representational theory of consciousness would now become, is what these abilities produce purely representational or do their products have nonrepresentational properties of some kind?

Step 3: Unity of Consciousness

In step 3, we turn to the first two properties of conscious systems as a whole mentioned above. Considered in the light of eliminativism and homuncular functionalism, the thing about conscious subjects that stands out most strikingly is that

1. each of us is precisely *one* of them, and

2. this one is aware at the same time, indeed in a single act of consciousness, of a great many items. A theory of the conscious subjects has to account for both the unitariness of the conscious subject and their consciousness of a great many things at the same time.

The latter is unity of consciousness. It comes in a number of kinds. In fact, mental unity in general comes in at least six different kinds. Four of them are kinds of unified consciousness:

2(i) unity of our cognitive elements (we can bring, for example, beliefs, desires, perceptions, intentions, and many other things to bear on a single situation);

2(ii) unified consciousness of our world (we are aware of a whole host of things around us in a single, unified representation) and

2(iii) unified consciousness of one's own representations;

2(iv) unified consciousness of self (one is aware of oneself as the "single, common subject" of one's experience, as Kant put it),

2(v) unified focus of a number of cognitive resources on a single item of attention; and

2(vi) unified behaviour (our behaviour is highly and multiply unified – think of a concert pianist playing a sonata).

Items (i) and (vi) have their own special features but there is a core of similarity among (ii) to (v), the kinds of unified *consciousness.* The core of the similarity is this:

> **Unity of consciousness:** A group of representations being related to one another such that to be conscious of any of them is to be conscious of others of them and of the group of them as a single group.

Given how central unified consciousness is to the conscious mind, it is remarkable how little attention it is has received in recent writings on consciousness, especially philosophical writings. Paul Churchland (1995, p. xx) includes it as one of the Magnificent Seven phenomena that a theory of consciousness has to account for and the notion is used by a few philosophers but in general the notion has received little attention. The unities of consciousness are the topic of Chapter 3.

That consciousness is unified has immediate implications for atomism and local realism. That consciousness is unified entails that consciousness is not a property that single representations or tiny groups of representations (e.g., a representation and a thought directed at it on the HOT model) could have by themselves. Nor it is something that could fruitfully be studied by studying single representations in isolation. At present, there is no theory of consciousness, representational or nonrepresentational, that provides an adequate account of the fact that consciousness is unified. To account for the various forms of unified consciousness within a purely *representational* theory might appear to be a tall order but it one that we will attempt to fill.

Step 4: The representational base of consciousness

As part of a theory of the conscious subject, we need a theory of the experiential requisites of something coming to appear to it. Since the theory of these experiential requisites that I will advance will play a number of important roles in the next chapter and beyond, I will lay it out here in a bit more detail than I am laying out the other aspects of my model of consciousness here.

What is the experiential basis of becoming conscious of how our acts of representing and what they represent appear to us? One standard view is that representations and/or their contents get to be conscious by becoming the object of another representational state, a thought (higher-order

thought (HOT) theory) or an experience (higher-order experience (HOE) theory). By contrast, I want to propose what I will call a single-order experience (SOE) theory. More specifically, in my view, it takes only one representation for a representation to become something of which we can be conscious. In fact, an act of representing can make us aware of three things. Consider the sentence:

1. I am looking at the words on the screen in front of me.

The representation this sentence expresses can make me aware of the words on the screen, obviously. But it can also make me aware of the representation, that the words are seen (rather than heard, imagined, touched). And it can make me aware of who is seeing the words, namely, me.

Let me introduce a term for this function. Let us call an act of representing that can make me aware of its object, itself and myself the *representational base* of my becoming aware of these items.

Representational base – an act of representing that can make me aware of its object, itself and myself

Almost any representation will do. Imagining Pegasus will do just as well as perceiving external objects such as computer screens. Indeed, representational states which have no apparent object such as pains or feelings of hunger do fine. Nor does a representation itself have to be recognized to provide a representational base for self-awareness. Just recognizing what is represented by a representation is enough for me to be aware that it is me who is aware of it.

That a single experience can serve as the representational base of consciousness of three different items, its object, itself and its bearer, does not require that conscious take the same form in all three cases. Some have argued that my 'consciousness' of my representations is actually an inference from the fact that I am representing something in them.[13] Others have argued (more plausibly in my view) that when one is conscious of one's own representations, this consciousness is not fully intentional, where what is meant by 'not fully intentional' here is that because the representation and what it makes us conscious of, namely itself in this case, are not separate, that representation is not straightforwardly *about* anything. It presents itself but it cannot appear in different ways in this presentation. If so, the usual means of misrepresenting are not present. And this suggests that intentionality of the usual sort is not present. Others have argued that when I am aware of myself as myself, I do not appear to myself *as* anything. We will deal with these matters in Chapter yy.

Indeed, not only can a representation present things, namely, itself and its bearer, without these things being its object, it can enter unified consciousness in a way other than becoming the object of such consciousness. As we will see early in Chapter 3, it is plausible to think of the relationship between individual representations and unified consciousness as a part/whole relationship, not a representation/object relationship.[14]

[13.] This is the transparency and displaced perception introduced in note 2.

[14.] This picture of the relationship between a representation and itself and a representation and unified consciousness is inconsistent with the forms of representationalist that assumes that representations connect to other things only by taking these other things as objects of representation. The picture is also probably inconsistent with HOT and HOE pictures.

Note that having the representational base for consciousness of a state is not the same thing as actually being conscious of it, nor indeed of myself. For example, in addition to a representation, a directing of attention or some cognitive apparatus might be needed to take advantage of the opportunity to represent and be aware of the representation. (Something like this might explain the point we saw Dennett trying to make that few if any species other than human beings have the ability to bring it about that a state becomes something it is like to have.) However, virtually any representation gives me all the *representation* I need. Frege and Husserl are said to have postulated a separate representation of self, an *Ich-Vorstellung.* On my theory, no such representation would be needed. Any garden variety representing of something in the world is enough for me to become aware of that representation and of myself as the thing that has it. [15]

Step 5:

The fifth step in the model of consciousness we aim to construct is easy to sketch, rather difficult to complete. In step 5 we aim simply to build a model that can deal with all the items in the list of requirements on a theory of consciousness given at the end of Step 1 which remain after steps 1 to 4 have been taken.

6. The Lie of the Land

I will close with a survey of the landscape of current philosophical work on consciousness in the analytic tradition. This work can be grouped into four broad categories. The second and third categories contain many variations:

1. What we call consciousness is really something else

On this view, consciousness will in the end be "analysed away" – what we misleadingly label 'consciousness' is really very different from anything that we pretheoretically take consciousness to be like. Perhaps the term 'consciousness' has been used for such a wide and diverse group of phenomena that it should be replaced with a number of more specific terms (P. S. Churchland 1983; we touched on this idea in Section 2). Or there is said to be nothing corresponding to our picture of consciousness as a highly integrated system able to unify many representations into one. Whatever,

[15.] This theory of self-awareness is remarkably powerful. In particular, it neatly avoids some of the problems that afflict other leading current theories of self-awareness, as we will see. As we will see, it is related to Peacocke's (2001) delta theory. In delta theory, having a property is sufficient in relevant cases for its subject to refer to and ascribe the property to itself. Interesting, Kant was already aware of the key elements of the representational base over two hundred years ago (Brook 2001).

the system in question is nothing like both common sense and the Cartesian and Kantian models of the mind picture it.

Contrary to a widely held view, right now there are no clear representatives of this eliminativism approach. Many will immediate ask about Paul and Patricia Churchland in this regard. Perhaps once upon a time they did flirt with eliminativism about consciousness. However, they have never advocated wholesale replacement of our consciousness talk as they did with respect to our intentional talk. In recent years they have backed away from whatever strength of eliminativism they originally held in connection with both. We will return to Paul Churchland shortly.

Some commentators will also think of Dennett's (1991) model of consciousness in this regard. I think that it would be a mistake to read Dennett this way. When he rejects what he calls the Cartesian Theatre picture of consciousness, he is rejecting a certain *theory* of consciousness. As he sees it, consciousness involves more interpretation by the cognitive system than has been thought, a system in turn that has less unity and cognitive structure than has been thought, and the resulting conscious states have less determinability and temporal stability than has been thought. None of this is to deny that something appropriately called consciousness exists. Indeed, Dennett insists that conscious states exist. He just want to deflate some philosophical pretensions about what consciousness is like (2000, pp. 369-70). However, the issue is a bit complicated. We will discuss it properly in the next chapter.

2. Consciousness is something very peculiar

This is the home of anti-representationalism and some related positions. On this set of views, consciousness (understood as qualia, the "felt quality" of representations) is held to be quite unlike anything else known in cognitive systems. Theorists in this group believe that consciousness could in principle be separated from other aspects of representations, in particular the aspects that causally affect other representations and behaviour (Nagel 1965, Block 1995, Chalmers 1996). One feature of this "neo-dualism", as Perry (2001) calls it, is heavy reliance on thought experiments about inverted spectra, inverted earth, zombies, colour-blind colour scientists, etc. These thought-experiments are designed to show that the content and causally effective aspects of a representation, on the one hand, and how it appears, on the other, could come apart. Support for this separation is often found in the explanatory gap mentioned earlier, our supposed inability to imagine how the felt qualities of experiences could connect to anything else in cognition or the brain (Levine 1983).

An important subclass of these theorists are those who argue that there is something central to consciousness that cannot be captured in a theory or description at all (McGinn 1991, Jackson 1986). Drawing on Chomsky's notion of cognitive closure, these theorists argue that something about consciousness is as closed to us as physics is to a field mouse (McGinn) or that what the felt quality of something is like cannot be captured in a theory about physical systems (Jackson) or that the nature of a point of view is systematically inexpressible in an objective theory done from an impersonal or God's-eye point of view (Nagel 1974. 19xx). Flanagan (1992) calls these people the new Mysterians. Not all mysterians are atomists. McGinn and Nagel are not, for example.

3. Consciousness is an aspect of individual representations or small groups of representations

We now consider theorists who are atomists but also representationalists. The core conviction of representationalists of whatever persuasion about atomism is that consciousness is simply a special form of representation or an aspect of some representations. Far from consciousness being some "layer" that could be peeled off a representation while leaving the rest of the representation intact, as in the inverted spectrum story, consciousness *simply is* representation: some kind or some aspect of ordinary representations. Whether it is consciousness of the world, consciousness of our representations, or consciousness of self, consciousness of something is nothing more than an aspect of representing that thing.

Many representationalists are atomists. These theorists view consciousness as a local property of certain psychological states. Among representational theories that hold that conscious states can be studied in isolation from conscious systems, HOT and HOE theories are a prominent subgroup. On this approach, a psychological state becomes conscious when a thought or other representation is directed at it. Thus, what makes something conscious is that it is the object of some other (itself nonconscious) thought or experience. The term 'higher-order' refers to the fact that the thought is about another psychological state, not some state of affairs in the world (Rosenthal 19xx; others xx).

Responding to the challenge of externalism, another prominent subgroup of atomistic representational theory tries to treat something appearing to be like something as an inference from representing something in the world around us. This group of theories urges that conscious states are "transparent": to the extent that our conscious states appear to us at all, they do not appear as anything different from the states of affairs of which they make us conscious. Here is the illustration we considered earlier. Return to the example of the square red patch. First focus on the patch. Now focus on your *awareness* of that square red patch Are you now aware of new properties? As we said, philosophers who answer 'no' include Harman (1990), Dretske (1995) Tye (1995) and Crane 19xx). As Dretske has put it, we are aware *through* our representations but we are not aware *of* our representations. Consciousness is said to be a displaced perception (Dretske 1995, Tye 1995), i.e., something that we infer from what representations do make us directly aware of, namely, states and events in the world, not something represented to us directly.

A third subgroup of atomistic, representational theories views conscious states as simply straightforward representations of a certain kind (Lycan 1987, 1996 – check; Flanagan 1992; Tye 1995). We have conscious representations of the world, we have conscious representations of some of our representations, and we have conscious representations of ourselves. On the one hand, we are directly aware of our representations, not just what they represent. We know more things and different things about them than we could gain by any process of inference. On the other hand, far from the "felt quality" and the functional properties of representations being separable, they are in fact one and the same thing. What makes these representations conscious? Opinions vary but one common answer is that we can describe them, report on them (Lycan 1987, 1996).

4. Consciousness is a property whole cognitive systems

Nonatomistic approaches to consciousness once dominated philosophy – think of Descartes and Kant. While this general approach has retained its popularity among cognitive psychologist, it has had fewer adherents among analytic philosophers, as we have seen. Nevertheless, it has not been abandoned entirely. Dennett (1991, 1996) and Paul Churchland (1995, 2002) are two important philosophers who continue the nonatomistic tradition. For Churchland, when the brain focuses attention on a representation, interprets it conceptually, holds it in short term memory, and updates its narrative of the "world-unfolding-in-time" (2001, p. xx), it is conscious of the representation. Others stress the fact that conscious states mean something to the cognitive system, are of use to it in a certain special way (Mack, URL). For Dennett, consciousness is a matter of one or more of the multiple drafts of various depictions in us achieving a certain kind of dominance in the dynamics of Pandemonium-architecture of competing information-parcels that make up cognition. We will examine these models in Chapter YY.

Most significantly, when researchers other than philosophers – psychologists, artificial intelligence researchers – work on consciousness, most of their work is nonatomistic and representationist, as we have noted. We will examine some of this work in Chapter yy, too. As will be clear by now, I favour some variant of the nonatomistic representational approach to consciousness.

A word about realism. Nonatomistic representationalists about consciousness fall into two subgroups, realists and nonrealists (or modified realists). All the thinkers that we have considered in this group so far are realists – they all hold that consciousness not only exists but exists as a property of at least human brains. There are a few philosophers whose view of consciousness, though not eliminativist, is so unlike the standard or even the Dennettian realist picture that they deserve their own subgroup. Rorty (19xx) and Davidson (19xx) come to mind. For Davidson, as we said earlier, consciousness arises out of a complex triangular interaction of oneself, other purposive beings, and the world and consciousness is, roughly, this triangulation generating stable attributions of consciousness. For Rorty, consciousness talk is simply an 'efficacious discourse', an interesting and useful way of talking about certain things that we take ourselves to be experiencing. I tend simply to assume that some form of realism is true, but I will say a bit about why this strikes me as a reasonable assumption near the end of the next chapter.

With this survey of the types of approaches to consciousness currently in play, we can now identify our strategy. First we will dispose of the approaches taken by groups (1) and (2). For this, all one has do is argue against anti-representationalism. If one can show that anti-representationalism is wrong, that would dispose of (1): if we can make a representational story fly, then there would be no motive to adopt eliminativism. Showing that anti-representationalism is wrong would also dispose of all variants of (2) (the latter variants because without anti-representationalism, there would be no motive to adopt mysterianism).That would leave the various views grouped in (3) and (4). Here the task will be to continue the argument we have started against atomism and local realism. If we do this successfully, then only (4), some version of nonatomistic representationalism, would be left. If we then dispatch non-or-modified-realist approaches, nonatomistic

representationalism realist approaches are all that would remain standing. Which would be fine with me.

In Chapter 2, we will examine anti-representationalism about how things appear to us by looking at Dennett's well-known attack on the central idea underlying the view, the idea of qualia. Chapter 3 shows, via a critique of HOT theory, that unified consciousness is what is needed to be conscious of particular representational states, not any kind of higher-order thought, and then lays out the various kinds of unified consciousness and explores their characteristics. Chapter 4 examine unified consciousness across time and its relationship (or lack of relationship) to being one and the same person over time. In Chapter 5, we turn to consciousness of self and examine some of the ways in consciousness of self is different from consciousness of other things. Chapter 6 will examine two recent attempts to connect consciousness of self to other things, Bermudez's attempt to show how it could develop out of simpler forms of consciousness and Cassam's attempt to argue that consciousness of self as subject of one's experiences requires consciousness of self as a thing in the world. [Campbell, too, maybe.] In Chapter 7, we will examine the challenge of externalism and cast a critical eye on Dretske's approach to meeting the threat to representationalism that he thinks externalism poses. Finally, Chapter 8 takes up the issue underlying everything else in the book, the issue of what a cognitive system that achieves consciousness would have to be like.

7. A Note on Method

[Maybe]

References

Baars, B. 1988. *A Cognitive Theory of Consciousness* Cambridge University Press

Block, N. 1995. On a confusion about a function of consciousness *Beh. and Brain Sci.* 18, 227-47

Brook, A. forthcoming a. "Judgments and Drafts Eight Years Later", in *Dennett: A Comprehensive Assessment*, ed. D. Ross. D. Thompson and A. Brook. (Cambridge, MA: MIT Press).

Brook, A. forthcoming b. "Consciousness and the Varieties of Externalism", *Supervenience, Causality, Mind and Action:* Proceedings (volume in preparation).

Brook, A. forthcoming c. "Kant. Self-Reference and Self-Awareness" and "Introduction", in *Self-Reference and Self-Awareness,* ed. Andrew Brook and Richard DeVidi. Amsterdam: John Benjamins.

Brook, A. forthcoming d. Review of Quassim Cassam, *Self and World* (Oxford, 1997), *Mind*

Brook, A. 1998. "Unified Consciousness and the Self" (Invited peer commentary on G. Strawson, "The Self") *Journal of Consciousness Studies* 5, 583-91

Brook, A. 1997. "Unity of Consciousness and Other Mental Unities", *Proceedings of the 19th Annual Conference of the Cognitive Science Society, 1997* (New York: Ablex Press), p. 875.

Brook, A. 1996. "Jackendoff and Consciousness", *Pragmatics and Cognition* 4, pp. 81-92.

Brook, A. 1994. *Kant and the Mind* New York: Cambridge University Press

Chalmers, D. 1996. *The Conscious Mind* Oxford: Oxford University Press

Churchland, P. M. 1995. *The Engine of Reason, the Seat of the Soul* MIT Press

Dennett, D. 1978a. Toward a cognitive theory of consciousness *Brainstorms* Bradford Books,149-73

Dennett, D. 1991. *Consciousness Explained.* Little, Brown

Dretske, F. 1995. *Naturalizing the Mind*. MIT Press

Flanagan, O. 1992. *Consciousness Reconsidered.* Cambridge, MA: MIT Press

Jackendoff, R. 1987. *Consciousness and the Computational Mind* MIT Press

Jackson, F. 1986. What Mary didn't know *J. Phil.* 83:5, 291-5

Kant, I. 1781/87. *Critique of Pure Reason*. Trans N. Kemp Smith. London: Macmillan 1927

Levine, J. 1983. Materialism and qualia: the explanatory gap *Pac. Phil. Quart.* 64, 364-61

Lycan, Wm. 1987. *Consciousness* MIT Press

Lycan, Wm. 1990. *Mind and Cognition.* Oxford: Blackwells Publisher

Lycan, Wm. 1996. *Consciousness and Experience* MIT Press

McGinn, C.1991. *The Problem of Consciousness.* Oxford: Blackwell Publishing Co.

Mack, A. http://psyche.cs.monash.edu.au/v7/psyche-7-16-mack.htm.

Nagel, T. 1965. Physicalism. *Phil. Rev.* 74, 339-56

Rosenthal, D. undated MS. State consciousness and higher-order thoughts

Tye, M. 1995. *Ten Problems of Consciousness.* Cambridge, MA: MIT Press