

Consciousness and the Varieties of Aboutness

MARTIN DAVIES

Thinking is special. There is nothing quite like it. Thinking – judging, believing and inferring – occurs in the natural order; but, at least sometimes, it seems hard to accept that there could be a fully satisfying reconstruction of thought in the terms favoured by the natural sciences – particularly, the physical and biological sciences (Davies, 1990). Some of our intuitions about thought are, in this way, similar to intuitions about consciousness; for consciousness, too, strikes many as somehow defying scientific explanation (McGinn, 1988). So, what is the connection between thought and consciousness? Is it, for example, only conscious beings that can be thinking beings?

States and events of thinking are semantically evaluable. Acts of judging, states of believing and other propositional attitude states, such as states of desiring, or intending, have semantic content or aboutness. In virtue of their semantic content, belief states, for example, can be evaluated as true or false, correct or incorrect, depending upon how the world turns out to be. The aboutness that is characteristic of the domain of thinking is often known as intentionality. So, what is the connection between intentionality and consciousness? Does intentionality require consciousness?

In his paper, ‘Consciousness, explanatory inversion, and cognitive science’ (1990a), John Searle delivers a strong affirmative answer to this question. According to Searle, it is not just that a thinking being needs to be a conscious being. Rather, a requirement of consciousness – of accessibility to consciousness ‘in principle’ – applies thought by thought, intentional state by intentional state. Thus, Searle’s *Connection Principle* (1990a, p. 586):

The ascription of an unconscious intentional phenomenon to a system implies that the phenomenon is in principle accessible to consciousness.

The Connection Principle plays a central role in Searle’s paper. Both the argumentative route leading up to it, and the route leading from the Connection Principle to the consequences that Searle draws, merit careful attention.

Contemporary cognitive science extends the notions of aboutness and semantic evaluability far beyond the domain of thinking, and far beyond what would ordinarily be regarded as the limits of accessibility to consciousness (even ‘in principle’). Searle regards this as a major error, and moves from the Connection Principle to dramatic conclusions concerning cognitive science (1990a, p. 589):

If we are looking for phenomena which are intrinsically intentional but inaccessible in principle to consciousness there is nothing there: no rule following, no mental information processing, no unconscious inferences, no mental models, no primal sketches, no $2\frac{1}{2}$ D images, no three-dimensional descriptions, no language of thought and no universal grammar.

The list includes many of the supposed glories of contemporary cognitive psychology and theoretical linguistics. So, it would seem that, if Searle is right, then the pretensions of those disciplines need to be re-evaluated quite radically.

These, then, are Searle's two aims: 'to show that we have no notion of an unconscious mental state except in terms of its accessibility to consciousness' – in short, to establish the Connection Principle – and 'to lay bare some of the implications of this thesis for the study of the mind' (1990a, p. 586). There is a vast amount to be said about each of these aims. In its original publication, Searle's paper (1990a) was accompanied by thirty-five commentaries, many – even most – of which were critical of one or another stage of his argument.

I agree with Searle that there is something special about intentionality, and that many of the states invoked in cognitive psychology and theoretical linguistics lack that special something. Indeed, the distinction to be recognised here is a deep one. We would fail to give the distinction its proper weight, if we said merely that intentional states and the states invoked in cognitive science are differing variations on the common theme of semantic evaluability. So far as his first aim goes, then, I am in substantial agreement with Searle. My disagreement with him is primarily over the *load-bearing* potential of the notion of accessibility to consciousness. I am not convinced that appealing to consciousness is the best way to make explicit what is distinctive about intentionality. Consciousness is the topic for Section 1.

But, even accepting that there is an important truth lying in the region of the Connection Principle, I disagree with Searle about the consequences of this truth for cognitive science. So far as his second aim goes, I claim that Searle's argument is undermined by the fact that he does not draw enough distinctions, at the outset, between different notions of aboutness. The varieties of aboutness provide the subject matter for Section 2. The remaining two sections focus upon the Connection Principle (Section 3) and upon cognitive science (Section 4).

1. Consciousness

Suppose that we were faced with the converse of the question to which the Connection Principle returns an answer. Suppose, that is to say, we were asked whether consciousness requires intentionality. Must a conscious being be a thinking being?

Confronted by this question – whether intentionality is a necessary condition for consciousness – we are liable to find ourselves with conflicting intuitions. On the one hand, to the extent that consciousness is just a matter of undergoing sensations and other experiences, it does not seem to require the cognitive achievements of judgement, belief and inference – the achievements of conceptualised thought. And this first intuition is strengthened further if there is reckoned to be an essential connection between thought and language; for it is natural to attribute experiences to infants and to other animals that lack language. On the other hand, to the extent that consciousness is a matter of a subject being aware of, able to think about and ultimately to report upon, his or her own mental states, then of course consciousness requires all that thinking requires.

A plausible explanation for this conflict of intuitions is that we are actually making use of at least two different notions – or families of notions – of consciousness. Kathleen Wilkes remarks (1988, p. 38) that, 'it is improbable that something bunching together pains, and

thoughts about mathematics, is going to be a reliable pointer to a legitimate natural kind', and Alan Allport is likewise sceptical that there is any such 'unitary phenomenon' as consciousness (1988, p. 162). These authors would, it is true, be almost as dubious about the idea that a mere binary distinction could bring order to this domain of enquiry. But let us, nevertheless, in a cautious and provisional spirit, begin by distinguishing between phenomenal consciousness and access consciousness (Block 1990, 1991, 1992, 1993, 1995; Davies and Humphreys, 1993).

1.1 Phenomenal consciousness and access consciousness: The initial distinction

The idea of *phenomenal consciousness* is the idea of 'something that it is like' to which Thomas Nagel directed our attention (1979, p. 166):

an organism has conscious mental states if and only if there is something that it is like to *be* that organism – something it is like *for* the organism.

We can say that a system is phenomenally conscious just in case there is something that it is like to be that system, and that a state of a system is a phenomenally conscious state if and only if there is something that it is like, for the system, to be in that state. Phenomenal consciousness, Ned Block tells us (1995, p. 228) is 'experience', and phenomenally conscious properties of states include 'the experiential properties of sensations, feelings and perceptions'. There is something that it is like to undergo sensations – to feel an itch, or a pain, or a tickle. There is also something that it is like to have perceptual experiences – to feel for the alarm clock when it rings, to recognise the face of a friend, to see the planet Venus in the evening sky, to see, hear, taste and feel cool running water. It is this notion – phenomenal consciousness – that is seen by Nagel, and by other philosophers such as Frank Jackson (1982, 1986) and Colin McGinn (1989), as constituting a locus of mysteriousness, and as creating an 'explanatory gap' (Levine, 1983, 1993) that cannot be bridged by the physical and biological sciences.

The idea of *access consciousness* is, to a very rough first approximation, the idea of availability for explicit verbal report (Fodor, 1983, p. 56). If a subject has a belief – say, the belief that the angle in a semicircle is a right angle – then typically she can verbally express the content of that belief (by saying, 'The angle in a semicircle is a right angle') and she can verbally report that she has that belief (by saying, 'I believe that the angle in a semicircle is a right angle'). If she is unable to do these things, then we may well say that if she has the belief at all then it is an unconscious or tacit belief.

Something similar goes for a sensation, like pain. If a subject has a pain then typically she can verbally express the pain (by saying, 'Ouch') and she can verbally report that she has that pain (by saying, 'I am in pain'). But there are also differences between the case of pain and the case of belief. Pain has something that belief lacks: it is not very plausible that there is anything that it is like, in Nagel's sense, to believe that the angle in a semicircle is a right angle. And belief has something that pain lacks: our subject can express the pain, but she cannot express the content of the pain, since pains plausibly do not have any semantic content. So, the notion of phenomenal consciousness applies only to the pain case, and one aspect of the notion of access consciousness applies only to the belief case. To the extent

that it separates the pain case and the belief case, the distinction between phenomenal consciousness and access consciousness achieves some credibility as a response to Wilkes's remark about 'bunching together pains, and thoughts about mathematics'.

1.2 Access consciousness: Some refinements

Access consciousness was introduced as availability for verbal report, but we need to refine that initial idea. Three different refinements will be considered.

Suppose that our subject verbally reports that she believes that the angle in a semicircle is not always a right angle. Suppose that she says this, not on any introspective grounds but, rather, on the authority of her psychoanalyst. Her analyst credits her with this belief by way of interpreting a range of inappropriate behaviour, and she trusts her analyst even though at this stage of the analytical process she is not yet able to identify with that belief – to recognise it 'from the inside'. Intuitively we do not want this kind of verbal report to be enough for access consciousness.

Similarly, suppose that our subject verbally reports that she tacitly knows that an anaphor is bound in its governing category while a pronominal is free in its governing category (Chomsky, 1986, p. 166). Suppose that she reports this on the authority of her linguistics professor, who has credited her with this piece of tacit knowledge by way of a partial explanation of her judgements about the grammaticality or otherwise of sentences such as:

I told them about each other.

I told them that Bill liked each other.

Then, once again, we would not want this to count as access consciousness of a state of tacit knowledge (cf. Searle, 1990b, p. 634).

There are important differences between attributions of unconscious knowledge and belief in the case of psychoanalysis and in the case of theoretical linguistics. But the common feature that matters here is that, in each case, the subject's ability to make her report depends upon much more than just her being in the state that is reported upon. In each case, she relies on the authority of a third party. We do not want these cases to count as examples of access consciousness. So, we should say that a state that has semantic content is access conscious if *simply in virtue of being in that state* the subject is able to express verbally the content of the state, and to report verbally that she is in the state. This takes us closer to a sufficient condition for an intuitively recognisable notion of access consciousness. But the requirement of verbal report raises a question as to whether we have formulated a necessary condition. In order to move to the first refinement of the idea of access consciousness, we need to replace or dilute the requirement of verbal report.

1.2.1 First refinement: Rational control of action

Allport, for example, having noted the consequences of the requirement of verbal report for a global aphasic (1988, p. 163), considers an alternative criterion of potential action (1988, p. 165):

In common usage, it seems, to be aware of something or conscious of something carries at least the implication that ‘something’ can guide or control my choice of *action*.

But, the problem that Allport then raises is that this criterion trades upon the distinction between voluntary and involuntary actions, and ‘there seems little distinction to be made between a “voluntary” action and one “consciously directed”’ (1988, p. 167). In short, the criterion seems to reintroduce the very notion that it is intended to clarify.

The criterion that Allport discusses is similar in spirit to Block’s (1995) account of access consciousness (p. 231):

A state is access-conscious if, in virtue of one’s having the state, a representation of its content is (1) inferentially promiscuous, i.e. poised to be used as a premise in reasoning, and (2) poised for *rational* control of action and (3) poised for rational control of speech. . . . I see [access consciousness] as a cluster concept, in which (3) – roughly, reportability – is the element in the cluster with the smallest weight, though (3) is often the best practical guide to [access consciousness].

We can take this – in which the verbal report requirement is heavily diluted – as our first refinement of the idea.

As the emphasis (in the original) indicates, the word ‘rational’ is not idle here. It is needed, rather as ‘voluntary’ is needed in the criterion discussed by Allport. In fact, whether or not there is a circularity in the offing, the appeal to reasoning and rational action in Block’s account of access consciousness suggests that this first refinement of the idea is pointing us towards the notion of a state that inhabits the domain of thinking – the space of reasons. For what seems to be important is that the state should have the kind of content that fits it for the domain of judgement, belief, inference, and intention – the conceptualised content of thoughts.

1.2.2 Second refinement: Thought content

On many conceptions of the relationship between thought and language, the verbal expression or report of a psychological state is a relatively contingent effect of something more fundamental, namely, a judgement, or more generally a piece of thinking. This reflection suggests a slightly different refinement of the idea of access consciousness, one that leads us even more rapidly to the notion of thought content. We might say that a state with semantic content is access conscious if, simply in virtue of being in that state, the subject is able to entertain in thought the content of the state.

So far as it applies to beliefs, there is something trivial about this refinement of the idea of access consciousness. Believing is a kind of thinking – along with framing a hypothesis, wondering whether it is so, doubting that it is so, wishing that it were so, and the like. So, of course, to be in a belief state is *ipso facto* to have the content of the belief available as a content of thought. But, suppose that there are also psychological states that have a different kind of semantic content from the content of thoughts. Suppose, in particular, that there are psychological states whose content is not conceptualised.

Thought content is a kind of conceptualised content: no one can think a thought with a particular content without possessing the constituent concepts of that thought. No one can believe that – or wonder whether, or doubt that, or wish that – the angle in a semicircle is a right angle, without possessing the concepts of angle, semicircle, right angle, and so on. In contrast, psychological states with non-conceptualised content would be contentful states that a subject could be in even though he or she did not possess the concepts that we would use to specify the states' contents. Then clearly, to be in such a psychological state would not *ipso facto* be to have the content of the state available as a content of thought.

Thus, we are led by this second refinement of the idea of access consciousness to the distinction between states – principally, propositional attitude states – whose content is necessarily conceptualised by the subject, and states that have semantic content even though it need not be conceptualised by the subject.

1.2.3 Third refinement: Higher-order thought

When we introduced the idea of access consciousness as availability for verbal report, we distinguished two components in the idea of verbal report. The subject can express verbally the content of the state, and can report verbally that she is in the state. When we strip away the requirement of verbal report, leaving just availability for thought, we can still retain two components. The first component now says that if a psychological state with semantic content is access conscious, then to be in the state is *ipso facto* to have the content of the state available as a content of thought. The second component says that if a psychological state is access conscious, then to be in the state is *ipso facto* to be in a position to judge that one is in that state. The refinement of the idea of access consciousness that we have just considered omits that second component. Let us now try adding it in.

The second component says that to be in an access conscious state is *ipso facto* to be in a position to judge that one is in that state. But there is a problem with that requirement. It is plausible that, for many psychological states, it is possible for a subject to be in the state without possessing the concept of that type of psychological state. Suppose, in particular, that it is possible to believe that the angle in a semicircle is a right angle without having the concept of belief. Then clearly, just to have the belief is not yet to be in a position to judge that one has the belief. If we are going to add in the second component, then we shall get closer to an intuitive notion of accessibility to consciousness if we restrict attention to subjects who do possess the concept of the type of psychological state in question.

Given that restriction, what difference does it make if we add in the second component? Of course, if a state is classified as not access conscious by the lights of the first component in the idea of access consciousness, then it is likewise classified as not access consciousness when we add in the further requirement of the second component. But, in fact, the second component itself suffices for the negative verdict in these cases.

In order to judge that she is in a state of a certain type and with a certain content, a subject needs to possess the concept of that type of state, and also to possess the concepts that figure in the specification of the content. Let us suppose that a subject can be in a state of tacitly knowing that an anaphor is bound in its governing category while a pronominal is free in its governing category, without possessing the concepts of anaphor, governing

category, and so on. Then to be in such a state of tacit knowledge is not *ipso facto* to be in a position to judge that one is in that state, even if one possesses the concept of tacit knowledge. This is simply because the content of the state is not *ipso facto* available as a content of thought. In effect, the second component contains the first component within it.

The difference that is made by adding in the second component becomes visible when we consider states with conceptualised content. If, for example, our subject has a belief then, trivially, the content of the belief is available as a content of thought. But intuitively, it is a further question whether the subject is, just in virtue of having the belief, in a position to judge that she has that belief. Indeed, intuitively this further question is very closely related to the question whether the belief is a conscious belief or not.

But, it is not clear that this third refinement of the idea of access consciousness quite captures the intuitive idea of a conscious belief. Suppose that the subject has a belief and is *ipso facto* in a position to judge that she has that belief. Then, it is very natural to suppose that there must be something about her belief state in virtue of which she is in a position to make that judgement. We might now ask which of two candidates it is that constitutes the belief's being a conscious belief. One candidate is the subject's being in a position to judge that she has that belief. The other candidate is the belief state's having the property – whatever property it is – that explains why the subject is placed in such a position. It is at least arguable that we should prefer the second candidate, rather than the first – the explanatory property, rather than the more dispositional property that it explains. But then, we must enter a reservation as to whether this notion of access consciousness – defined in terms of the more dispositional property – goes quite to the heart of the notion of a conscious belief.

Indeed, following this line of thought, we can even suggest that access consciousness, as it is defined in this third refinement, is not necessary for consciousness, as it applies to beliefs. For if what matters for conscious belief is the explanatory property, then we can envisage a belief state having that property, and so being a conscious belief, even though the subject lacks the concept of belief, so that this notion of access consciousness is not even applicable (Peacocke, 1992, pp. 152–3).

It should be acknowledged, though, that more work needs to be done if this appeal to an explanatory property of belief states is to be ultimately satisfying. A subject who is in pain is, we suppose, in a position to judge that she is in pain, provided only that she possesses the concept of pain. In this case, we may say that it is the pain state's being a phenomenally conscious state that explains why being in the state disposes the subject to judge that she is in pain. But, having denied that there is anything that it is like to believe, for example, that the angle in a semicircle is a right angle, we cannot tell a parallel story about the grounds of the subject's dispositions to make judgements about her own beliefs.

However, whatever the complexities of that issue, there are other grounds for denying that this notion of access consciousness goes to the heart of consciousness as it applies to belief states. For access consciousness does not seem to be sufficient for consciousness as it applies to beliefs. There are imaginable cases in which a subject has a belief, and is thereby in a position to judge that she has that belief, and even does judge that she has that belief, but where both the first-order and the second-order beliefs would be intuitively reckoned as

unconscious beliefs. An example could be based on unconscious guilt about an unconscious belief, since the most intelligible basis for the guilt would be recognition – also unconscious, of course – that one has that belief (cf. Peacocke, 1992, p. 154).

Adding the requirement that the subject be in a position to judge that she has a certain belief takes us close to an intuitive notion of consciousness – of conscious belief. But we have found reasons to doubt whether access consciousness, thus refined, is either necessary or sufficient for consciousness on that intuitive conception. To the extent that we doubt whether the third refinement of the idea of access consciousness captures the idea of a conscious belief, we shall also doubt that the so-called higher-order thought theories of consciousness are adequate. These are theories of consciousness that equate a psychological state's being a conscious state with the subject's actually having a thought about the state, to the effect that she is in that state (Nelkin, 1986, 1989a, 1989b; Rosenthal, 1986, 1993). For exactly similar queries can be raised about both the necessity and the sufficiency of the higher-order thought theorist's conditions for being a conscious belief.

1.3 Does consciousness require intentionality?

After drawing a distinction between phenomenal consciousness and access consciousness, we have been considering three ways of refining the initial idea of access consciousness. The first and second refinements both point in the direction of conceptualised content as being the key notion. The third refinement goes further, but does not quite capture the intuitive notion of a conscious belief. These distinctions – introduced in a cautious and provisional spirit – may help us better to understand what is at issue in Searle's arguments. They may suggest a doubt, for example, as to whether the notion of accessibility to consciousness can bear the argumentative weight that Searle places upon it. We shall take up these matters in Section 3.

Now, however, we can return to the question with which we began this section. Does consciousness require intentionality? If we consider access consciousness first, then the close tie with intentionality is immediately obvious. The first and second refinements point directly towards the notion of intentionality, and the third refinement clearly requires intentionality since the subject is said to be in a position to make a judgement, and so arrive at a belief, about her own mental state. Even when the mental state under consideration does not itself exhibit intentionality – a pain, for example – if the state is access conscious in virtue of the fact that the subject of the state, possessing the concept of pain, is in a position to judge that she is in pain, then the subject is capable of mental states that do have intentionality.

If we turn now to phenomenal consciousness, the plausible answer to the question – at least initially – is surely that phenomenal consciousness does not require intentionality. Suppose that a phenomenally conscious mental state – a pain – is also access conscious, in the sense that the subject of the pain, possessing the concept of pain, is able to judge that she is in pain. Then, as we have already noted, we may say that it is the pain state's being a phenomenally conscious state that explains the subject's disposition to judge that she is in pain. The pain's being a phenomenally conscious state is not itself dependent upon the subject's ability to make that judgement.

However, we should note two dissenting voices here. First, higher-order thought theories of consciousness treat phenomenal consciousness as a kind of access consciousness. So they return the verdict, across the board, that consciousness requires intentionality. We have, though, already noted reasons for querying the account of consciousness that these theories offer.

The other dissenting voice comes from Gareth Evans (1982). In a recommendation that is fundamental to his programme, Evans says that we should (1982, p. 123):

take the notion of *being in an informational state with such-and-such a content* as a primitive notion for philosophy, rather than . . . attempt to characterize it in terms of belief.

He then deploys this notion of a state of information in his account of perceptual experience (1982, pp. 226–7):

In general, we may regard a perceptual experience as an informational state of the subject: it has a certain *content* – the world is represented a certain way – and hence it permits of a non-derivative classification as *true* or *false*. . . .

The informational states which a subject acquires through perception are *non-conceptual*, or *non-conceptualized*.

However, in Evans’s view, states of information are also implicated in unconscious processes that take place in our brains, and in the explanation of the correct ‘guesses’ of blindsight subjects, for example. So, what makes the difference between those informational states that constitute conscious perceptual experiences and mere unconscious informational states? Evans’s answer is that (1982, p. 158):

we arrive at conscious perceptual experience when sensory input is not only connected to behavioural dispositions . . . but also serves as the input to a *thinking, concept-applying, and reasoning system*.

It is only when we have a subject whose conceptualised judgements are appropriately sensitive to the non-conceptualised information content of perceptual states that we have a subject for whom those perceptual states amount to conscious experiences.

There is an important difference between this account and that of the higher-order thought theorists. The judgements mentioned in Evans’s account are judgements about the world, rather than about mental states. But still, we might say that Evans treats phenomenal consciousness as something like access consciousness. Certainly for Evans, as for the higher-order thought theorists, consciousness – the phenomenal consciousness of perceptual experience – requires intentionality. An experiencing subject is also a thinking subject.

This is not the place to assess Evans’s view about the link between consciousness and intentionality. But we should notice that his appeal to the notion of information, and to non-conceptualised content more generally, does not depend upon that link. (We shall mention Evans’s view again in Section 2.4, below.)

2. The Varieties of Aboutness

Someone might point out that Searle’s claim in the Connection Principle (1990a, p. 586):

The ascription of an unconscious intentional phenomenon to a system implies that the phenomenon is in principle accessible to consciousness.

is subject to either a weaker or a stronger construal, and then seek to make peace between Searle and cognitive science by imposing just the weaker construal upon his words. On the weaker construal, Searle would merely be saying that psychological states that are inaccessible to consciousness have an importantly different kind of aboutness, or have aboutness in a different way, from conscious – or potentially conscious – psychological states, such as beliefs. On the stronger construal, he would be saying that the only genuine aboutness belongs to (at least potentially) conscious mental states; so that supposedly psychological states that are inaccessible to consciousness do not have genuine aboutness at all. On the weaker construal he would be allowing the possibility of a variety of grounds for non-derivative and non-trivial semantic evaluability, while on the stronger construal he would be ruling out that possibility.

Unfortunately for the peacemakers, Searle appears to intend the stronger construal of his conclusions. He does draw a distinction between genuine (or ‘intrinsic’) intentionality and mere ‘as if’ intentionality (1990a, p. 586); and he allows that the unconscious states invoked in cognitive science participate in ‘as if’ intentionality. But this does not permit any kind of vindication of cognitive psychology and linguistics, since ‘as if’ intentionality is far too inclusive (*ibid.*):

I can, for example, say of my lawn that it is thirsty, just as I can say of myself that I am thirsty. But it is obvious that my lawn has no mental states whatever. When I say that it is thirsty, this is simply a metaphorical way of describing its capacity to absorb water. . . . We say such things as, ‘The carburetor of the car *knows* how rich to make the mixture,’; ‘The thermostat on the wall *perceives* the changes in the temperature,’; ‘The calculator *follows rules* of arithmetic when it does addition,’ and ‘The Little Engine That Could is *trying very hard* to make it up the mountain.’ None of these attributions is meant to be taken literally, however.

And, likewise in an earlier presentation of the position (1989, p. 198):

water flowing downhill behaves *as if* it had intentionality. It *tries* to get to the bottom of the hill by ingeniously *seeking* the line of least resistance, it does *information processing* in order to *calculate* the size of rocks, the angle of the slope, the pull of gravity, etc.

If, as the Connection Principle states, genuine intentionality requires accessibility to consciousness and ‘as if’ intentionality is the only alternative, then trivialisation is the fate that awaits great tracts of psychology and linguistics as those disciplines are currently conceived.

Still, the peacemaker’s idea is not without relevance to the assessment of Searle’s arguments. For it may be that, although Searle intends the stronger construal of his conclusions, it is only the weaker construal that is strictly licensed by the arguments that he gives. So, in the remainder of this section, I shall note – in as neutral a way as possible – a variety of kinds of aboutness, none of which appears to be as trivial as mere ‘as if’ intentionality. We begin with genuine, intrinsic intentionality.

2.1 *Attitude aboutness*

First, then, our belief states have aboutness. Suppose that Fiona believes that Venus is a planet. Her belief is about an object, Venus, and a property, being a planet: it is the belief that the former exemplifies the latter. More generally, beliefs, desires, hopes, fears, wishes, and intentions – the propositional attitudes – have aboutness. We shall reserve the term ‘intentionality’ for this kind of aboutness: *attitude aboutness*.

We have already noted that attitude aboutness involves conceptualisation. For present purposes, we can think of possessing a concept as having a particular cognitive ability; in the most basic cases, this is the ability to think of an object or a property in a certain way. Possessing a concept involves appreciating certain special, or canonical, rational connections between thoughts involving that concept and other thoughts or experiences. Thus, to pursue our well-worn example, Fiona might have a way of thinking of the planet Venus which we would describe as thinking of Venus as the first heavenly body to appear in the evening (whether or not Fiona could articulate that specification herself). Part of what is involved in that cognitive ability is appreciating that evidence that is made available by Venus’s appearing in the evening sky is directly relevant to thoughts in which Venus is thought about in that way.

Clearly, Fiona could have other ways of thinking of the same planet Venus – for example, a way that would privilege evidence made available by Venus’s appearance in the morning sky. In this neo-Fregean framework (Evans, 1982; Peacocke, 1986, 1992) ways of thinking, and equivalently concepts, are discriminated extremely finely – as finely as is required by differences in the cognitive significance of thoughts in which the same object or property is thought about. So, if Fiona has two thoughts which are about the same planet, to the effect that it exemplifies the same property, yet these thoughts differ in their cognitive significance for her, then the thoughts involve two different ways of thinking (of either the planet or the property).

Fine-grainedness is one characteristic of conceptualisation, and so of intentionality, or attitude aboutness. Another characteristic is that the cognitive abilities that constitute possession of concepts, and are deployed together in one thought, can be recombined in countless other ways, enabling the thinking subject to entertain countless other thoughts. This is the characteristic of conceptualisation that Evans spells out in the Generality Constraint (1982, p. 104). A third characteristic is that possession of any concepts of objects and properties at all seems to require possession of a range of spatial concepts and concepts for kinds of medium sized material bodies, along with some appreciation of the fact, for example, that one way for a body to exist unperceived by a particular subject is for the body to exist at a place other than that at which the subject is located (Strawson, 1959; Evans, 1980).

On this kind of account, conceptualisation – and so, intentionality – is a highly non-trivial cognitive achievement. Thinking is special.

2.2 *Linguistic aboutness*

Propositional attitude states are not the only things that have aboutness. A second clear example is provided by sentences – or utterances of sentences – of a public language. The

English sentence ‘Venus is a planet’, like Fiona’s belief, is about Venus and the property of being a planet. Also, like Fiona’s belief, the sentence is true – or correct – if and only if Venus is indeed a planet. But arguably, the meanings of words in a public language cut less finely than ways of thinking do. Suppose that W is a word whose semantic value (reference) is an object or property V. Then it may be that two language users each count as understanding W as having the meaning that it does actually have, even though they think about V – the semantic value of W – in two quite different ways.

Belief states and utterances of sentences both have aboutness. But it would be wrong to assume that a philosophical account of the aboutness – or meaning – of sentences will take just the same form as an account of the aboutness – or intentionality – of beliefs. It is plausible that the meaning of public language sentences has something to do with the way that utterances of those sentences are conventionally used to communicate particular messages – to pass on beliefs from speaker to hearer. But that kind of account clearly cannot be applied to beliefs themselves. On the contrary, that kind of account of linguistic meaning seems to take the aboutness of beliefs and other propositional attitudes for granted. As Searle says (1983, p. 5), ‘the direction of logical analysis is to explain language in terms of Intentionality’; and (1983, p. 27):

An utterance can have Intentionality, just as a belief has Intentionality, but whereas the Intentionality of the belief is *intrinsic* the Intentionality of the utterance is *derived*. . . .

The mind imposes Intentionality on entities that are not intrinsically Intentional by intentionally conferring the conditions of satisfaction of the expressed psychological state upon the external physical entity.

So, as well as the *attitude aboutness* of beliefs and desires, we have the *linguistic aboutness* of sentences and utterances. To the extent that conventional signs, and perhaps also maps and pictures, inherit their aboutness from the intentionality of beliefs in a similar way, we might also call this *conventional aboutness*. But, whatever we call it, we should note that the aboutness of language is derived from the aboutness of propositional attitudes. It is a kind of derived intentionality: derived, rather than intrinsic, but still genuine, rather than merely ‘as if’. Whatever we are doing when we say that a sentence or utterance has meaning or aboutness, we are not engaging in a pretence that the linguistic item is a bearer of mental states.

2.3 Indicator Aboutness

At the beginning of Paul Grice’s paper ‘Meaning’ (1957), we find the example, ‘Those spots mean (meant) measles’, along with the remark (1989, p. 213):

I cannot say, ‘Those spots meant measles, but he hadn’t got measles’ . . . That is to say, in cases like the above, *x meant that p* and *x means that p* entail *p*.

Grice called this notion of meaning *natural meaning*, as opposed to non-natural meaning which is a matter of the speaker’s intentions.

Other examples of this kind of meaning – phenomena that have meaning in the sense that they *indicate* something about the world – are not difficult to find. Thus, we say, ‘Those

spots mean – or indicate – measles’, ‘Those clouds mean – or indicate – rain’, and ‘The existence of thirty rings in that tree trunk means – or indicates – that the tree was thirty years old when it was cut down’.

This *indicator aboutness* is clearly distinct from the other two kinds that we have listed so far. As Grice points out, we cannot consistently say, ‘Those clouds mean – or indicate – that it will rain; but in fact it will not rain’. But we can consistently say, ‘The content of Fiona’s belief is that it will rain; but in fact it will not rain’. And we can say, ‘The meaning of (the utterance of) that sentence is that it will rain – that was the message that was communicated; but in fact it will not rain’. Indicator aboutness, at least in its simplest form, differs from attitude aboutness and linguistic aboutness in that it does not allow for the possibility of falsehood or misrepresentation.

There are broadly two views that a theorist might have concerning this kind of aboutness. On the one hand, a theorist might treat indicator aboutness as a further species of derived intentionality. The idea would be that the aboutness of the clouds – the clouds’ meaning that it will rain – is inherited from the attitude aboutness of the belief that someone could form on the basis of observing the clouds. On the other hand, a theorist might regard indicator aboutness as a feature of the world that is not logically dependent upon the propositional attitudes of observers. On this second view, the fact that clouds in general, and those clouds in particular, mean that it will rain is something to be discovered, and then to be relied upon in forming beliefs. But the fact obtains, whether it is discovered or not.

As between these views, we can say that the second is more plausible as an explication of what Grice was pointing to with the notion of natural meaning. The first view makes indicator aboutness a highly relative notion, since different observers will arrive at different beliefs, depending upon their different bodies of background knowledge. There is such a notion – those clouds mean rain to me, but not to her – but it seems to be distinct from the notion of natural meaning. And the first view does not make it obvious why indicator aboutness closes off the possibility of falsehood. An aboutness that is derived from the intentionality of beliefs should allow for the same possibilities of error as does attitude aboutness.

The second view is to be preferred, then. As Fred Dretske says (1986, p. 18): ‘Naturally occurring signs mean something, and they do so without any assistance from us.’ And we can take it that indicator aboutness is to be explicated in terms of reliable causal covariation between events of two types. In the example of the clouds, the two types would be occurrences of a certain kind of cloud formation and occurrences of rain shortly afterwards.

Given this background, what should we say about the familiar example of the fuel gauge: ‘The position of the fuel gauge means – or indicates – that the tank is almost empty’? Since the fuel gauge is an artefact, surely – it might seem – our intentions and beliefs enter into its aboutness. But it is important to distinguish indicator aboutness from a different notion here. If the states of the fuel gauge do reliably covary with the states of the fuel tank, then the position of the fuel gauge, towards the bottom of the scale, means or indicates that the tank is nearly empty – and it does so without any assistance from us. Suppose that the fuel gauge starts to malfunction, the covariation becomes unreliable, and the needle takes up a position towards the bottom of the scale even when the tank is full. Then the position of the fuel

gauge no longer indicates that the tank is almost empty, even though that is what it is supposed to indicate, and what the designer intended it to indicate. Upon this way of carving up the territory – in which, to be sure, there may be an element of stipulation – indicator aboutness is independent of the propositional attitudes of either observer or designer. But the function of the fuel gauge – what its purpose is, what it is for – is in this case a mind-dependent matter.

How does indicator aboutness fit into Searle's scheme? On the view that we are taking, indicator aboutness is not a kind of derived intentionality. So, according to Searle's taxonomy, it is either genuine, intrinsic intentionality, or else merely 'as if' intentionality. Since indicator aboutness extends to spots, clouds, tree rings, and fuel gauges – things that have no mental life and are not conscious – indicator aboutness is not generally accessible to consciousness, even in principle. So, according to the Connection Principle, that would rule out genuine, intrinsic intentionality. Is indicator aboutness just 'as if' intentionality, then?

Searle certainly includes a room thermostat – which is relevantly similar to a fuel gauge – amongst his examples of 'as if' intentionality. But the classification is not entirely satisfying. I may say that it is as if my lawn is thirsty, as if the lawn wants it to rain, or even that it is as if the lawn believes that it is going to rain soon. ('Brave lawn, refusing to die, and waiting for the rain that it believes will arrive shortly.')

But these metaphorical attributions are not constrained by any requirement of a reliable causal covariation, as attributions of indicator aboutness are.

On the other hand, a metaphorical attribution of intentionality carries with it a metaphorical attribution of all that is required for intentionality. So, it is said to be as if the lawn possesses concepts, as if the lawn has fine grained ways of thinking about objects and properties, as if the lawn measures up to the Generality Constraint, and as if the lawn appreciates that an object may exist unperceived by the lawn in virtue of existing at a place different from where the lawn is located. It is not, of course, *very much* as if all these things are true; but then, some metaphors are better than others. However, there is nothing corresponding to any of this in the conditions for attribution of indicator aboutness. Indicator aboutness is not fine grained, it is not subject to the Generality Constraint, and it does not carry with it any appreciation of objects existing unperceived.

The possible objection to Searle's scheme of things is thus that it involves the unsatisfying classification of central cases of indicator aboutness as examples of mere 'as if' intentionality on a par with the attribution of thirst, desire, and belief to a lawn.

A possible response is to say that nothing very much hangs on this classification since the notion of indicator aboutness can be of only limited theoretical interest. Two reasons could be given for not attaching great importance to indicator aboutness. The first is simply that it does not allow for the possibility of falsehood. The second is that, if indicator aboutness is just a matter of causal covariation, then it is too cheap: too many things causally covary with each other.

This response has some force. But it still leaves a potential problem for Searle. For it is at least open that there may be a notion of aboutness that is just as independent of our intentions and beliefs as indicator aboutness, but which lacks the two features mentioned in the response. Any notion of aboutness which, like Dretske's functional meaning (1986,

p. 22), combines the idea of causal covariation with that of natural function (cashed out in terms of evolutionary selection) will be a candidate.

However, rather than take that line of thought any further now, let us turn briefly to two further kinds of aboutness.

2.4 Experiential aboutness

Perceptual experiences present the world as being one way or another. They present objects as having certain properties. Thus, we may hear a sound as coming from the left; or we may see a box as being cubic and about four feet in front of us. This *experiential aboutness* is closely related to attitude aboutness; for we can form beliefs on the basis of experiences such as these. Thus we may have an auditory experience and come to believe that the sound is coming from the left; or we can have a visual experience and come to believe that the box is cubic and about four feet in front of us.

But still, it is worth distinguishing the content of experiences from the content of the beliefs that we form on the basis of those experiences. For, it is an important point about belief content that, in order to believe that a box is cubic, one has to possess the concept of a cube. But, it is arguable that merely having a box presented in experience as cubic does not, in the same way, require possession of that concept.

As we noted at the end of Section 1, this is Evans's (1982) view. Perceptual experiences have a kind of aboutness that does not require conceptualisation by the subject of the experiences. Furthermore, Evans recommended treating experiences as informational states, where the notion of such states is taken as a primitive and not explained in terms of belief. On this view, then, experiential aboutness is not a species of derived intentionality.

Experiential aboutness is clearly distinct from the other three notions that we have introduced. It is different from attitude aboutness in not requiring conceptualisation. It is different from linguistic aboutness in not being derived. And it is different from indicator aboutness in permitting falsehood: perceptual experiences can present the world as being different from the way that it really is.

2.5 Subdoxastic aboutness

Finally, there are unconscious psychological states that have aboutness. At least, such states are invoked throughout cognitive psychology and theoretical linguistics. In these disciplines, the unconscious processes that lead up to experience – and ultimately to belief – are certainly described in content-using, or semantic, terms. The very idea of information-processing psychology is committed to there being information in the visual system or the auditory system, for example: information about other states of the creature and also about states of the external world. The states that carry this information are often said to be mental representations.

Let us call this unconscious psychological aboutness *subdoxastic aboutness*, in line with Stephen Stich's (1978, p. 499) labelling of states which 'play a role in the proximate causal history of beliefs, though they are not beliefs themselves' as subdoxastic states. It is this notion of aboutness that is principally *sub judice* in the context of Searle's arguments leading up to, and onwards from, the Connection Principle.

But to the extent that this notion is legitimate at all, it is distinct from attitude aboutness, linguistic aboutness, and indicator aboutness for the same reasons that experiential aboutness is different from those other three. Subdoxastic aboutness is distinct from attitude aboutness since, like experiential aboutness, it is a kind of non-conceptualised content. (Indeed, it is between those two kinds of non-conceptualised content that Evans (1982, p. 158) distinguishes in terms of serving ‘as the input to a thinking, concept-applying, and reasoning system’.) Subdoxastic aboutness is distinct from linguistic aboutness since it is not derived. Subdoxastic aboutness is also distinct from indicator aboutness, since it allows for the possibility of misrepresentation. We can say, for example, that a state of the auditory processing system represents the presence of a sound coming from the left even though there is not in fact any sound coming from the left. And finally, subdoxastic aboutness is quite unlike experiential aboutness – the fourth variety of aboutness – since it is not tied to consciousness.

2.6 Simplifications

Having distinguished five kinds of aboutness, we can simplify the discussion to follow by moving the notions of linguistic aboutness and experiential aboutness away from centre stage. In the case of linguistic aboutness, the reason is that there is nothing that is in dispute concerning this notion. In the case of experiential aboutness, it is not quite clear whether this notion is legitimate from Searle’s point of view. Since it is the aboutness of perceptual experiences, it is closely tied to consciousness, so that is a point in its favour. On the other hand, Searle has not expressed any sympathy for the idea of non-conceptualised content. In any case, the issues are too complex to pursue here.

Then, we can achieve some further simplification by noting that, while the notions of indicator aboutness and subdoxastic aboutness are distinct, still a substantive theory of subdoxastic aboutness might well draw upon and refine the notion of indicator aboutness. In particular, it may be that a theory of subdoxastic aboutness will make use of the idea of covariation on both the input side – causal antecedents – and the output side – causal consequences – along with the idea of natural function. It is plausible that, from these resources, we can construct a notion that allows for the possibility of misrepresentation, and is rather more expensive than indicator aboutness.

Classification of core cases of subdoxastic aboutness as examples of mere ‘as if’ intentionality will then be just as unsatisfying as the similar classification of cases of indicator aboutness. On the one hand, metaphorical attributions of ‘as if’ intentionality are not constrained by requirements of causal covariation or of natural function. On the other hand, there is nothing in attributions of subdoxastic aboutness that corresponds to the heavy commitments of conceptualisation. Subdoxastic aboutness presents itself, then, as a *prima facie* candidate for being a ground of non-derivative and non-trivial semantic evaluability, distinct from genuine, intrinsic intentionality.

With our attention focused now upon attitude aboutness and subdoxastic aboutness – upon genuine, intrinsic intentionality and the kind of semantic content that is widely invoked in cognitive science – we can consider Searle’s arguments.

3. The Connection Principle

The Connection Principle plays a central role in Searle's argument. His twin aims are to establish the principle and then to draw out its consequences for the study of mind. One way to avoid the alleged consequences of the Connection Principle – 'no rule following, no mental information processing, no unconscious inferences, no . . .' – and defend cognitive science would be totally to reject the Connection Principle itself. Indeed, this would be the natural response from a friend of cognitive science who shared Searle's background assumption that aboutness is a roughly unitary phenomenon – that there is only one kind of non-derived, non-trivial aboutness. This kind of friend of cognitive science is likely to regard accessibility to consciousness as an optional extra, distinguishing a subclass of semantically evaluable states with no scientific integrity – 'a scattered and probably uninteresting subpart of the full cognitive structure', as Chomsky puts it (1976, p. 163).

The position that I favour is intermediate between Searle and Chomsky. Thinking is special – as Searle says. Cognitive science is in good order – as Chomsky says. Crucially, I disagree with Searle about the Connection Principle's consequences for cognitive science. The first of three claims that I shall make in this section is that, if the notion of subdoxastic aboutness – as distinct from attitude aboutness – is so much as allowed to get to first base, the Connection Principle cannot then deliver a threat to its legitimacy.

On the other hand, I agree with Searle that there is something special about attitude aboutness, though I disagree with him about the *load-bearing* potential of the notion of accessibility to consciousness. The second claim in this section is that what is special about thinking is not best captured by the Connection Principle's criterion of accessibility to consciousness in principle.

The third claim in this section is simply that the argument for the Connection Principle is far from watertight. As Searle himself remarks (1990b, p. 634):

Many commentators point out . . . that the argument does not demonstrate the Connection Principle with absolute certainty. They are right about that. The argument is explanatory though not demonstrative . . .

To deny that the argument is watertight is not, however, to deny that there is an important truth lying in the region of the Connection Principle. Indeed, my view is that there is something true and important there, something that is not yet well understood.

3.1 *The consequences of the Connection Principle*

Let us allow, for the moment, that attitude aboutness – genuine, intrinsic intentionality – requires accessibility to consciousness, just as the Connection Principle says. And let us accept that many of the semantically evaluable states invoked in cognitive psychology and theoretical linguistics are not, in the relevant sense, accessible to consciousness. Still, this does nothing to threaten the legitimacy of the notion of subdoxastic aboutness. Certainly it does not suggest that the only thing to be said about subdoxastic aboutness is that it is mere 'as if' intentionality.

The first step in the argument for the Connection Principle is that ‘There is a distinction between intrinsic and as-if intentionality’ (1990a, p. 586). The reason for insisting upon this distinction is that (1989, p. 198):

the price of giving it up would be that everything becomes mental, because relative to some purpose or other anything can be treated *as if* it were mental.

We shall scarcely dispute this, any more than we should dispute the distinction between dogs and ‘as if’ dogs. The price of giving up that distinction would be that everything would be counted as a dog, since anything at all can be treated as if it were a dog. But it would be a mistake to move from the importance of this distinction to the idea that all that can relevantly be said about a cat is that it is an ‘as if’ dog. Similarly, it would be a mistake to think that all that can relevantly be said about a state that lacks genuine, intrinsic intentionality is that it is a case of ‘as if’ intentionality.

The friend of cognitive science wants to describe states as being non-derivatively and non-trivially semantically evaluable – as being assessable as correct or incorrect, true or false. In opposition, Searle argues that this talk of semantic evaluation should be replaced by talk about hardware and about function (1990a, p. 591). But, in order not to be question begging, Searle’s argument must not simply build in an assumption saying that attitude aboutness is the only possible source of non-derived, non-trivial semantic evaluability. So, let us imagine that it is explicit that we start by allowing for the possibility of other kinds of aboutness.

It should be clear that, if that possibility is allowed at the beginning, the later stages of the argument cannot then close it off. Suppose that the argument establishes that attitude aboutness requires accessibility to consciousness. Then, whatever is in principle inaccessible to consciousness is not attitude aboutness. But it does not follow that it is not a source of semantic evaluability. The situation would remain the same if the Connection Principle were replaced by a principle linking intentionality to something other than accessibility to consciousness. If attitude aboutness requires X, and the states invoked in cognitive science lack X, then those states do not have attitude aboutness. But they may yet be semantically evaluable in virtue of having some other kind of aboutness.

There is a response that might be made, on Searle’s side of the disagreement. It might be said that, while a state that is inaccessible to consciousness could be semantically evaluable, this would only be in some extended and trivial sense, since it would be a case of mere ‘as if’ intentionality. But there is a fallacy that must be avoided here. It is true that the state in question would lie within the domain of ‘as if’ intentionality, along with states of dry lawns and water flowing downhill. But the state in question may have genuine semantic properties of its own – semantic properties that are lacked by the states of lawns and streams. The state may have genuine aboutness; not genuine attitude aboutness or intentionality, of course, but genuine aboutness all the same.

It appears, then, that the argument for the Connection Principle could only have consequences threatening to cognitive science if it were to be augmented by the assumption that the only distinctions that are worth drawing, amongst kinds of aboutness, are those between genuine, intrinsic intentionality, derived intentionality, and ‘as if’ intentionality. But, so far, we have no good justification for that assumption. What is needed, if Searle’s

second aim is to be carried out, is an independent reason to cast doubt on the notion of subdoxastic aboutness and its cognitive science kin. We shall return to this point in Section 4.

3.2 *The criterion in the Connection Principle*

A theorist who maintains, against Searle, that there is a *prima facie* viable and explanatory notion of subdoxastic aboutness might still accept the Connection Principle as imposing a requirement upon attitude aboutness. For this theorist can allow for – may even insist upon – the importance of the difference between attitude aboutness and subdoxastic aboutness. This theorist is free to embrace the specialness of thinking – of concept possession, judgement, belief and inference – and might offer accessibility to consciousness as the distinctive mark of the domain of thinking. However, my own view is that what is special about thinking is not best captured by the Connection Principle’s criterion of accessibility to consciousness in principle.

In ‘Tacit knowledge and subdoxastic states’ (1989a), I addressed the question whether subdoxastic states – such as states of tacit knowledge – are like or unlike propositional attitude states. Chomsky says (1986, p. 269):

Suppose that the facts were different, and that we could become conscious, by thought and introspection, that we do in fact make use of these rules and principles in our mental computations. Then, I think, one would have no hesitation in saying that we know them. If so, then cognizing would appear to have the properties of knowledge in the ordinary sense of the term, apart, perhaps, from accessibility to consciousness.

There are two claims in this strand of Chomsky’s thought. One is that the difference between subdoxastic states and states of ordinary knowledge and belief is to be characterised in terms of accessibility to consciousness. The other is that the distinction is of no great importance for serious explanatory purposes.

I suggested, in contrast, that three intuitive differences between beliefs and subdoxastic states – accessibility to consciousness, inferential integration, and conceptualisation – add up to a *prima facie* case for a principled distinction. But, I also suggested that it is difficult to found the distinction upon the difference between accessibility and inaccessibility to consciousness, given our limited understanding of that notion.

One reason that I offered began from the fact that we do make use of the idea of unconscious beliefs (1989a, p. 136). These states would have to be allowed to count as accessible, if the criterion were to be plausible. One way – the only obvious way – of protecting the idea that these belief states are accessible to consciousness is to introduce into the criterion the idea of being accessible save for the presence of a blocking mechanism. But then this strategy must itself be restricted, lest it be used to classify even paradigm cases of subdoxastic states as being accessible if only there were not a blocking mechanism. And whatever restriction we impose is likely to go more nearly to the heart of the distinction between beliefs and subdoxastic states than does the idea of accessibility to consciousness.

Searle allows that unconscious beliefs are accessible in principle (1990a, p. 586): ‘there is nothing in principle inaccessible to consciousness about the Freudian unconscious’, and in

response to an objection, he links the notion of accessibility in principle with that of a blocking mechanism (1990a, p. 595). It is predictable, then, that many of his commentators (e.g. Block, 1990; Chomsky, 1990; Clark, 1990) object that relaxing the requirement of accessibility to accessibility ‘in principle’ risks obliterating the contrast between accessibility and inaccessibility altogether.

Another reason for not relying upon the notion of accessibility to consciousness to provide a criterion for the distinction between beliefs and subdoxastic states is that the notion is simply unclear. We can add some clarity by using the notions of phenomenal consciousness and access consciousness (cf. Block, 1990). But still we seem to be pointed elsewhere for the fundamental ground of the distinction. On the one hand, phenomenal consciousness is of no direct help here, for even if some kind of experience were to be associated with the occurrence of a state that had hitherto been classified as subdoxastic, that would not be enough to turn it into a belief state (1989a, p. 138). Access consciousness, on the other hand, might seem more promising. Its title suggests as much. But we need to recall the various refinements of the idea.

The first two refinements that we considered (in Section 1.2) point more or less directly in the direction of conceptualised content as providing a more fundamental criterion. We noted, for example, that Block’s account of access consciousness (1995, p. 231) suggests that what is important is that a state should have the kind of semantic content that fits it for the space of reasons, the domain of judgement, belief, inference, and intention. As Block points out in his commentary on Searle (Block, 1990, p. 597), if there is not a more fundamental criterion in the offing, then it is difficult to see what interest could attach to the question whether a state is well situated with respect to reasoning and reporting processes.

The third refinement of the idea of access consciousness goes beyond the notion of conceptualised content by introducing the further requirement that the subject should be in a position to judge that she is in the state in question. But, as we noted, the effect of this requirement is only visible when it is used to distinguish between conscious and unconscious beliefs. And this means that the added requirement is no help for our present purposes, since accessibility in principle is supposed to include unconscious beliefs. Exactly the same reflections apply, of course, if we replace the subject’s being in a position to judge that she has a certain belief by the belief state’s having the property that explains why the subject is placed in such a position.

It seems, then, that the key to the distinction between beliefs and subdoxastic states is to be found, not in the notion of accessibility to consciousness, but in the idea of conceptualisation. The content of beliefs and other propositional attitude states is necessarily conceptualised by the subject of those states, while the content of subdoxastic states does not require conceptualisation by the subject of those states.

We can relate this proposal to the previous subsection, and to the next one. First, a principle linking intentionality and conceptualisation has no dire consequences for cognitive science. If attitude aboutness requires conceptualisation, and the states invoked in cognitive science have non-conceptualised content, then those states do not have attitude aboutness. But they may yet have some other kind of aboutness that constitutes a source of non-derived and non-trivial semantic evaluability. Second, even if accessibility to consciousness in

principle is not a load-bearing notion, there may still be an important link between attitude aboutness and consciousness, mediated by the requirement of conceptualisation. Indeed, conceptualisation is a close relative of the notion that plays a pivotal role in Searle's argument for the Connection Principle, the notion of aspectual shape.

3.3 Aspectual shape and the Connection Principle

Searle's argument for the Connection Principle turns upon the claim (Step 2 of his argument) that 'Intrinsic intentional states . . . always have aspectual shapes' (1990a, p. 587). The notion of aspectual shape is explained (ibid.):

Whenever we perceive anything or think about anything, it is always under some aspects and not others that we perceive or think about that thing.

and then further elucidated by way of some examples (ibid.):

When you see a car it is not simply a matter of an object being registered by your perceptual apparatus; rather you actually have the conscious experience of the object from a certain point of view and with certain features.

A man may believe . . . that the star in the sky is the Morning Star without believing that it is the Evening Star.

A man may . . . want to drink a glass of water without wanting to drink a glass of H₂O.

What both the explanation and the examples suggest is that Searle's notion of aspectual shape is much the same as the Fregean notion of a mode of presentation. In the domain of judging, believing and inferring, objects and properties are always thought about under a mode of presentation – they are thought about in one way rather than another. And perceptual experiences make possible thoughts about objects and properties under perceptual demonstrative modes of presentation (Evans, 1982, Chapter 6; Peacocke, 1983, Chapters 5 and 6). So, the doctrine about aspectual shape is something to which we are committed, if we think that conceptualisation is what is distinctive of attitude aboutness – or genuine, intrinsic intentionality – and if we have a neo-Fregean view of concepts. Let us agree: intentionality requires a Fregean sense-reference distinction.

The question now is whether there is a close link between conceptualisation and consciousness – a link that would sustain the Searlean conclusion that a conceptualised state must be accessible to consciousness at least in principle. Although Searle himself does not make use of the distinction, we can consider this question for access consciousness and then for phenomenal consciousness.

The answer to the question whether there is a close link between conceptualisation and access consciousness is that the connection is all too close. Searle's notion of accessibility to consciousness in principle is intended to apply to unconscious beliefs, and we have just noted (in Section 3.2) that the notions of access consciousness that meet that requirement (the first and second refinements) are scarcely distinguishable from the notion of conceptualisation. Interpreted in terms of access consciousness, the Connection Principle says that genuine, intrinsic intentionality involves a special kind of content – conceptualised

content. By my lights, this is a correct principle. But it is surely less than Searle was seeking to establish.

The question whether there is a close link between conceptualisation and phenomenal consciousness is much more difficult to answer. But Searle's argument – from pivotal point to conclusion – makes good sense when construed as directed towards a positive answer to this question. When Searle says (Step 3 of his argument: 1990a, p. 587):

The aspectual feature cannot be exhaustively or completely characterized solely in terms of third person, behavioral, or even neurophysiological predicates. None of these is sufficient to give an exhaustive account of aspectual shape.

it is impossible not to be reminded of Nagel on the elusiveness of phenomenal consciousness (Nagel, 1979, p. 167):

If physicalism is to be defended, the phenomenological features must themselves be given a physical account. But when we examine their subjective character it seems that such a result is impossible. The reason is that every subjective phenomenon is essentially connected with a single point of view, and it seems inevitable that an objective, physical theory will abandon that point of view.

There is a complication to be dealt with in this construal of Searle, since we have said that there is nothing that it is like to have the belief, for example, that the angle in a semicircle is a right angle. Phenomenal consciousness belongs to sensations and perceptual experiences (Section 1.1). But we can take it that Searle is using a notion of consciousness that is tied to Nagel's idea of 'something that it is like', but extends to what we would intuitively regard as conscious beliefs.

Against that background, suppose that we were to grant both that genuine, intrinsic intentionality involves aspectual shape (Step 2), and that the fundamental philosophical account of aspectual shape adverts to what it is like to be in the intentional state in question (something very like Step 3). Then it would be natural to argue from those two premises to the conclusion that there is something problematic about the notion of an intentional state for which there is nothing that it is like to be in that state. This is essentially the point Searle reaches at his Stage 4: 'Now we seem to have a contradiction' (Searle, 1990a, p. 588). Furthermore, it would then be practically inevitable to suggest a particular way of resolving the apparent contradiction. A state to which the philosophically fundamental account of aspectual shape cannot apply directly could be credited with aspectual shape derivatively from some other state to which it stands in an appropriately intimate relationship. One candidate for the required relationship would be causal antecedence. Thus, Searle (1990a, p. 588):

[T]he only fact about the neurophysiological structures [realising states that there is nothing that it is like to be in] that corresponds to the ascription of intrinsic aspectual shape is the fact that the system has the causal capacity to produce conscious states and processes where those specific aspectual shapes are manifest.

There would be a legitimate query to be raised as to whether the relationship of causal antecedence is quite intimate enough. After all, subdoxastic states – which are precisely not states with genuine, intrinsic intentionality – are defined by Stich (1978, p. 499) as states

that ‘play a role in the proximate causal history of beliefs’. But this is a matter of detail. It is clear enough how to proceed from premises corresponding to Steps 2 and 3 of Searle’s argument to something very close to his conclusion.

Modulo the extension of the notion of phenomenal consciousness to encompass conscious beliefs, Searle’s actual argument for the Connection Principle makes good sense when it is construed as an argument for a link between intentionality and phenomenal consciousness. But that is not yet to say that the argument is compelling.

The reconstructed argument begins from two premises. One premise is that genuine, intrinsic intentionality involves aspectual shape. To the extent that aspectual shape is equivalent to Fregean mode of presentation or sense, this premise is something to which we are committed by what we have said about conceptualisation. The other premise is that the fundamental philosophical account of aspectual shape adverts to what it is like to be in the intentional state in question. This premise is controversial if aspectual shape is taken to be equivalent to Fregean mode of presentation or sense. There are at least two reasons for this. First, this premise depends crucially upon the extension of the notion of phenomenal consciousness from sensations and perceptual experiences to conscious beliefs. But that extension was entertained only in order to construe Searle’s argument, and not as a positive recommendation. Second, this premise depends upon an extrapolation to all modes of presentation from the case of conscious perceptions: ‘Aspectual shape is most obvious in the case of conscious perceptions’ (Searle, 1990a, p. 587). This premise can, of course, be rendered uncontroversial by a stipulation about the way in which the notion of aspectual shape is to be taken. But then, the controversy is simply shifted to the other premise. For nothing that we have said about conceptualisation leads to the view that there is always something that it is like to be in an intentional state.

As we are construing Searle’s argument for the Connection Principle, then, the problem that it faces is this. The argument turns upon the claim that ‘Intrinsic intentional states . . . always have aspectual shapes’. If a requirement of phenomenal consciousness is built into the notion of aspectual shape, then the argument onwards from the pivotal point is plausible, but the claim itself is controversial. If a requirement of phenomenal consciousness is not built into the notion of aspectual shape, then the claim itself is plausible, but the onward argument limps. (Some such dilemma as this seems to motivate a number of Searle’s commentators: e.g. Lloyd, 1990; Rosenthal, 1990.)

But even if – as Searle himself accepts (1990b, p. 634) – the argument for the Connection Principle is not absolutely compelling, still there may well be an important truth connecting intentionality and phenomenal consciousness. Genuine, intrinsic intentionality – attitude aboutness – involves conceptualisation, and conceptualisation involves senses or modes of presentation. Amongst modes of presentation, those demonstrative modes of presentation that are afforded by perceptual experience constitute particularly clear examples. Suppose now that we could argue that some theoretical primacy attaches to perceptual demonstrative modes of presentation. Suppose, even, that we could argue that in order to be able to think about objects at all, a subject needs to be able to think about objects under perceptual demonstrative modes of presentation. Then there would be a deep connection between

intentionality and consciousness, just as Searle says, although not one that holds intentional state by intentional state.

Whether or not we can establish any theoretical primacy for perceptual demonstrative thoughts is not a topic for this paper. Considerations in favour of such primacy would begin from the role of demonstrative thoughts in the explanation of a subject's actions upon objects in her environment (Perry, 1979). But for now, it is enough to raise the question and to note that the issue is far from being well understood.

4. Cognitive Science

We have seen (Section 3.1) that an attack on the legitimacy of subdoxastic aboutness must draw upon resources going beyond the Connection Principle. One possible line of objection to subdoxastic states, tacit knowledge of rules, information processing, and the rest of cognitive science's core notions, is generated by a threat of trivialisation. Certainly, the need to avoid trivialisation constitutes a challenge for any development of the notion of subdoxastic aboutness.

Concerning 'as if' intentionality, Searle says (1990a, p. 587):

Everything in the universe follows laws of nature, and for that reason everything behaves with a certain degree of regularity, and for that reason everything behaves *as if* it were following a rule, . . . For example, suppose I drop a stone. The stone . . . follows the rule $S = \frac{1}{2}gt^2$.

So, perhaps Searle would claim that, once we move away from conscious, conceptualised attitude aboutness, there is no way to prevent the notion of tacit knowledge of a rule from becoming utterly trivial, so that a wooden block sliding on a smooth surface would be credited with tacit knowledge of the rule ' $a = F/m$ ', for example. Certainly the block's motion conforms to that rule.

An initial reply to this worry is to say that, as the notion is used in information-processing psychology, a system that possesses knowledge of a rule is a system that has a resource enabling it to perform inference-like transitions between input and output states that are themselves states with aboutness – states that have semantic content. Knowledge of the rule ' $a = F/m$ ', for example, would permit the transition from an input state with the semantic content that the force is n units to an output state with the semantic content that the acceleration is n/m units. If we are asking about the presence or absence of that piece of knowledge, then the relevant input-output transitions are not from an input state that *is* a force to an output state that *is* an acceleration, but from an input state that *represents* a force to an output state that *represents* an acceleration. Likewise with Searle's example. Trivialisation only threatens if we cannot block the attribution of tacit knowledge of the rule ' $S = \frac{1}{2}gt^2$ ' to the falling stone. And, on the face of it, we can block that attribution by pointing out that there are no plausible candidates for states of the stone that represent the elapsed time, or the distance travelled.

In fact, the notion of tacit knowledge that I sketched in 'Tacit knowledge and subdoxastic states' (1989a) is even further removed from the trivialisation that might be suggested by Searle's remark. For attributions of tacit knowledge are not licensed simply by an input-

output relation; they require, in addition, a particular causal-explanatory structure in the mechanisms that mediate between input representations and output representations (Evans, 1981; Davies, 1986, 1987, 1989b, 1995).

It would be entirely fair to respond, on Searle's behalf, that this way of avoiding trivialisation in the notion of tacit knowledge of rules is only as convincing as the idea that we can assign aboutness to the input and output states of these processing mechanisms. If the only notion of aboutness that is applicable to those states is 'as if' intentionality, then we have scarcely advanced the matter at all. But we should need to be given some good reason to suppose that this is so, in the context of our earlier discussion of indicator aboutness and subdoxastic aboutness, and particularly given the differences that we noted between those notions and mere 'as if' intentionality.

If the notion of tacit knowledge of rules as I have sketched it here, and developed it elsewhere, is at all close to the notion actually used in cognitive science, then Searle underdescribes the practice of cognitive science when he says (1990a, p. 593):

A typical strategy in cognitive science has been to try to discover complex patterns such as those found in perception or language and then to postulate combinations of mental representations which will explain the pattern in the appropriate way. . . . Epistemically, the existence of the patterns is taken as evidence for the existence of the representations.

We can see this underdescription in Searle's use of the examples of language acquisition and the vestibulo-ocular reflex (VOR).

4.1 Language acquisition

In the case of language acquisition, the starting point of Chomsky's theory is that the state attained by a normal adult language user involves a body of knowledge that is brought to bear – in ways that are not specified by Chomsky – in both perception and production of sentences. A theory of how the knowledge is used is a theory of performance; a theory of the body of knowledge itself is a theory of competence. Given this starting point, there is then an argument for postulating a substantial innate endowment that is specific to the task of acquiring language.

In early versions of the theory – the Standard Theory (Chomsky, 1965) – the attained body of knowledge is supposed to consist primarily of a set of rules (a grammar). The innate endowment is then supposed to include at least two components. There is a body of knowledge about the universal features of humanly possible grammars. And there is an evaluation procedure for selecting one grammar from amongst those that are both attainable and consistent with the linguistic data that is available to the developing child.

In more recent versions – the Government-Binding Theory (Chomsky, 1986) – the attained body of knowledge and the innate endowment are more closely related. Now, the innate endowment, or initial state of the language faculty, is a body of knowledge that can be specified by a set of principles. In these principles, the values of certain parameters have yet to be fixed – they are like switches waiting to be set. The attained state is a body of knowledge specified by those same principles, but now with the values of the parameters set

– in each case to one of finitely many possible values. The role of the linguistic data available to the child is simply to enable the parameters to be set to one of their possible combinations of positions.

There are powerful reasons for the changes in the theory over a period of twenty five years or so, but the whole story about language acquisition lapses if the starting point is rejected; that is, if it is denied that the attained state is a state of knowledge, a state with semantic content. The issue between Searle and the explanatory practice of theoretical linguistics would be seen most clearly if Searle were to reject this starting point; and, indeed, it would be natural for him to do so. After all – at least so far as the attained state involves knowledge of principles – the principles are typically quite inaccessible to consciousness. But in fact, Searle focuses the dispute on the innate endowment: knowledge of universal grammar.

The reason that he gives for not focusing on the attained state is that (1990a, p. 593) ‘grammars of particular languages, like French and English, whatever else they contain, obviously contain a large number of rules that are accessible to consciousness’. But this is not clearly a very good reason. At least three considerations are relevant here.

First, when Searle speaks of the rules of grammars of French or English, he is not using the term ‘grammar’ as it is used within the theory under discussion. Searle is using the term as roughly equivalent to what Chomsky calls ‘pedagogical grammar’ (1986, p. 6):

a full list of exceptions (irregular verbs, etc.), paradigms and examples of regular constructions, and observations at various levels of detail and generality about the form and meaning of expressions.

But this is quite different from the body of knowledge that is reckoned to be the attained state of the language faculty.

Second, within Chomsky’s theory there is a distinction between core grammar and periphery (1986, p. 147). The distinction is, as Chomsky remarks, theory internal: the differentiation within the attained state is dictated by the collection of principles that are included in the initial state of the language faculty. The core grammar is just that part of the attained grammar that is constituted by the collection of universal principles with their parameters now set. The periphery is ‘whatever is added on in the system actually represented in the mind/brain of the speaker-hearer’ (ibid.). Now, it may be that some of what is in the periphery is accessible to consciousness. But, even so, the core grammar would have provided an example of aboutness without accessibility to consciousness – an example upon which to focus the issue.

Third, we do, of course, have conscious access to many of the consequences of the body of knowledge in the attained state, and particularly to many of the consequences of core grammar. Thus, most English speakers are aware that there is ‘something wrong’ with the sentence:

John and Peter wanted me to vote for each other

even though it is clear how that sentence would have to be interpreted (as saying that John wanted me to vote for Peter and Peter wanted me to vote for John). But conscious access to

consequences of the principles is not the same as conscious access to the principles themselves.

Concerning the attained state – particularly, core grammar as that is conceived in the principles and parameters framework – Searle could have made two claims exactly parallel to those that he makes about universal grammar. The first claim is that, if the very idea of a body of knowledge that is inaccessible to consciousness is illegitimate and incoherent – if the hypothesis that there is such a body of knowledge is ‘empty’ (1990a, p. 593) – then it is futile to try to defend the idea on the grounds that it is invoked in the ‘best theory’. Searle’s second claim is this. If we have to live without the idea of unconscious knowledge, and to view the hypothesis that there is such a body of knowledge as empty, then that is really no great loss. For, Searle says, whatever evidence was supposed to favour the hypothesis can be ‘much more simply accounted for’ by alternative hypotheses about hardware and function (1990a, p. 594).

It is in this second claim that the underdescription of cognitive scientific practice can be seen. For the alternative hypotheses scarcely go beyond the evidence plus the assumption that the brain is somehow implicated in the causal history of that evidence. Thus, James Higginbotham says (1990, p. 609):

No hint of theory is in sight here; rather, we have a replacement of explanations of the customary sort with sheer description of the facts that they were supposed to explain. It is thus highly misleading for Searle to say that the evidence for universal grammar is ‘more simply’ accounted for in his view, since the alleged account consists merely in stating the empirical facts.

Having to do without the idea of unconscious bodies of knowledge really would mean a diminution in explanatory power.

To the extent that Searle does not see this, it is presumably because he is already convinced that the notion of tacit knowledge of rules is trivial. If the friend of tacit knowledge cannot prevent the attribution of tacit knowledge of the rule ‘ $a = F/m$ ’ to a body sliding on a smooth surface, then he is indeed open to an objection of the form that Searle envisages. For, whatever the attribution of tacit knowledge is supposed to explain can be ‘much more simply accounted for’ by the alternative (hardware) hypothesis that, for a body of a given mass, the acceleration produced is proportional to the net force applied. But what we have already argued is that the threat of trivialisation can be blocked.

The second claim is not, then, a telling point against the practice of theoretical linguistics. The first claim, in contrast, is surely correct. If a notion is illegitimate, and hypotheses using the notion are empty, then it is pointless to insist that the theory in which these hypotheses figure is the best theory. But no dire consequences follow, unless it is shown that the notion of a body of knowledge that is inaccessible to consciousness *is* an illegitimate notion. It is not shown to be illegitimate by the argument for the Connection Principle, provided that we distinguish the aboutness of the body of knowledge from attitude aboutness. Nor has it been shown to be a trivial notion, yielding hypotheses that are devoid of empirical content.

4.2 *The vestibulo-ocular reflex*

The vestibulo-ocular reflex (VOR) generates eye movements that compensate for head movements in such a way as to keep stable the image of the world (or the image of a moving object that is being tracked) on the surface of the retina. It is a feedforward, rather than a feedback, system; the VOR does not operate by detecting visual errors. Indeed, (Churchland and Sejnowski, 1992, p. 356):

The basic VOR circuit consists of detection [of head movements] by transducers, projection to the vestibular nucleus in the brain stem, and projection from there to cranial nerve nuclei, where motor neurons originate that project to the eyes muscles.

As a result of these very direct linkages, the VOR operates much more rapidly than the optokinetic system – which uses feedback from slippage of the retinal image – and the pursuit system – which is used for tracking small objects. Indeed, the VOR operates at extraordinarily high speeds: typically the compensating eye movements begin around 12-14 milliseconds after the onset of the head movement. By way of comparison, eye movements that are produced by the optokinetic and pursuit systems are delayed by more than 80 milliseconds (Miles and Lisberger, 1981, pp. 275–6; Churchland and Sejnowski, 1992, p. 353).

One way of describing the VOR, then, is simply as a system by means of which head movements cause equal and opposite eye movements. We might call this a mechanical input-output description. Another way to describe the VOR is as a system in which certain information processing takes place. The starting point for this second kind of description is that transitions take place, not just from head movements of certain velocities to eye movements of certain velocities, but from representations of head movement velocities to representations of eye movement velocities (Churchland and Sejnowski, 1992, p. 357–8):

If we assume that neurons projecting to the vestibular nuclei in the brainstem carry signals specifying head velocity, and that motor neurons carry signals specifying muscular contractions to produce an eye velocity, then the computational action, so to speak, is located between these two at the several vestibular nuclei, lumped together for convenience as ‘VN’.

It is only against the background of this second kind of description that there is any question of crediting the system with tacit knowledge of rules relating head velocity and eye velocity. To invoke knowledge of such rules given only the first kind of description – the mechanical input-output description – would be to trivialise the notion. It would be like crediting a wooden block sliding on a smooth surface with knowledge of the rule ‘ $a = F/m$ ’.

So, the assumption of input and output states that have semantic content is a necessary condition for a description of the intervening causal transitions in terms of knowledge of rules. But, it is not yet a sufficient condition for crediting the system with tacit knowledge of a rule or generalisation relating head velocity to eye velocity. According to the notion of knowledge of rules that I favour, tacit knowledge of a rule in a processing system requires that the various input-output transitions that are in conformity with the rule should have the same causal explanation – roughly, should be mediated by the same component mechanism

within the system (see again Evans, 1981; Davies, 1986, 1987, 1989b, 1995). And that is not a completely trivial requirement.

A properly constrained information-processing account of the operation of the VOR in terms of ‘a deep unconscious rule’ cannot, then, simply be replaced by a hardware account as Searle suggests (1990a, p. 591): ‘When I look at an object while my head is moving, the hardware mechanism of the VOR moves my eyeballs.’ There would be a loss of explanatory power in that replacement. Someone who denies that there is any such loss is likely to be overlooking the necessary starting point for an information-processing account, and supposing instead that the account in terms of rules is set against the background of only the mechanical input-output description of the reflex.

The computational problem posed by the VOR is, in any case, rather more complex than the starting point for an information-processing account might suggest. For, as Churchland and Sejnowski go on to explain (1992, p. 358):

VN is actually an area of convergence, receiving not only vestibular signals, but also signals from smooth-pursuit neurons and saccadic-burst neurons. . . . The computational problem for the VN centers on what transformation should be applied to the input vector carrying the three kinds of information so as to solve the problem of how the eyes should move.

Searle appears to reject the whole idea of an exploration of the VOR in information-processing – or computational – terms, preferring some neurophysiological elaboration of this sketch (1990a, p. 591):

What actually happens is that fluid movements in the semicircular canals of the inner ear trigger a sequence of neuron firings that enter the brain over the eighth cranial nerve. These signals follow two parallel pathways . . . in the brain stem and cerebellum and they transform the initial input signals to provide motor output ‘commands’, via motoneurons that connect to the eye muscles and cause eyeball movements.

But against this we have to set the fact that, while there will surely be a neural account of the VOR, this is no reason to exclude other levels of explanation unless those other levels have been shown to be inevitably headed in the direction of triviality or incoherence. So, we should pay some heed to such a statement as this, from David Robinson (1981, p. 463):

When we ask how the brain works, the question is perceived quite differently by people working at the many levels of the nervous system. Certainly it is necessary to know how the hardware of the nervous system works . . . but we all recognize that the solutions to these problems are not an end in themselves. They are the means that will enable us to understand how the brain processes information . . .

Robinson’s own work on the VOR (e.g. Robinson, 1981, 1987; Cannon, Robinson and Shamma, 1983; Cannon and Robinson, 1985) contains a fine example of the role of an information-processing description in neuroscientific research.

Robinson’s research suggests that, in general, the firing rate, R , of the ocular motor neurons (responsible for initiating the muscle contractions that produce eye movement) is characterised by an equation:

$$R = b + kE + r(dE/dt)$$

where b is a constant – the background discharge rate of the neuron – and the other two terms are proportional to the eye displacement, E , and the eye velocity, dE/dt , respectively (1981, p. 467). Given that some muscle force is required just to maintain the eye in a position displaced from straight ahead, even if the eye is not moving, the presence of the kE term in the equation makes good sense. But the problem that it poses is that, while conversion of head velocity information to an eye velocity command is a simple – even trivial – computational matter, obtaining information about eye displacement from information about velocity requires integration. (This is not integration in the sense of bringing information together – as in cross-modal integration – but integration as in high school calculus.)

The project of devising a model of a neural integrator provides a classic example of interaction between levels of explanation. Very roughly, the idea that a neural integrator is needed comes from a high level task analysis (Robinson, 1987, p. 1915):

The need for a neural integrator was apparent as soon as anyone trained in systems analysis sat down and looked at this reflex.

An initial computational – in fact, neural network – model (Cannon, Robinson and Shamma, 1983) had some attractive features, and solved the problem of how to avoid integrating the background firing rate along with the velocity signal. But the units in the neural network model differed in important ways from the corresponding population of real neurons, believed to be in the vestibular nuclei. So, considerations from neurophysiology motivated a revised model, returning the research to the computational level (Cannon and Robinson, 1985).

Whether or not Robinson's story about the neural integrator is ultimately correct, the investigation – moving as it does amongst several levels of description – constitutes an episode of strikingly explanatory research in cognitive science. There is nothing in Searle's arguments that reveals a hitherto concealed vacuity in this research programme.

To be sure, there will be theorists who regard the use of a level of description at which neural firing rates are credited with semantic content as a level having primarily heuristic importance – guiding the development of detailed neurophysiological theories. These theorists will go beyond mere appreciation of inter-level interaction in the direction of more reductive aspirations, and to that extent they will regard descriptions in terms of aboutness, information, and rules as having diminishing importance as science progresses. But even these reductionist theorists can still agree that the descriptions in terms of aboutness are far from trivial. They are certainly constrained more tightly than the mere 'as if' description of my lawn as thirsty, or as desiring or even believing that it will rain. So, even a reductionist theorist can reject Searle's claim that, in an information-processing account of the VOR, 'All the intentional ascriptions are *as-if*' (1990a, p. 591).

5. Conclusion

The first two sections of this paper were taken up with some distinctions. In Section 1, there was the distinction between two notions of consciousness – phenomenal consciousness and access consciousness – along with three refinements of the second of those notions. In

Section 2, there were distinctions between five different notions of aboutness that might plausibly find a place in our thinking about psychology and the mind. We then simplified the discussion by focusing on just two of those notions – attitude aboutness and subdoxastic aboutness.

With these distinctions in place, we have been concerned, in the last two sections, with Searle's two aims. He aims to establish the Connection Principle, and then to show that it has serious negative consequences – 'no . . . , no . . . , no . . .' – for cognitive science.

In Section 3, I argued for three claims of my own. First (Section 3.1), the Connection Principle does not, by itself, threaten the legitimacy of the notion of subdoxastic aboutness. If it appears to do so, then this is because a question-begging assumption has been used. This is the assumption that the only distinctions worth drawing in the area of aboutness are the distinctions amongst genuine, intrinsic intentionality, derived intentionality, and 'as if' intentionality. Second (Section 3.2), what is special about thinking is best captured by a criterion of conceptualisation rather than a criterion of accessibility to consciousness. Third (Section 3.3), although the argument for the Connection Principle is far from watertight, there may well be a link between intentionality and consciousness, mediated by the requirement of conceptualisation and the idea of perceptual demonstrative modes of presentation. More work is needed to improve our understanding here.

Although the Connection Principle does not pose an independent threat to subdoxastic aboutness, information processing, knowledge of rules, and the like, those notions do face a challenge. They must avoid the threat of trivialisation. In Section 4, I argued that this challenge can be met. Indeed, the appearance of a serious threat depends, to some extent, upon an underdescription of the practice of cognitive science.

Thinking – concept possession, judgement, belief, desire, inference, intention – is special. There is nothing quite like it. Certainly, unconscious information processing and tacit knowledge of rules are not quite like it. But those explanatory resources deployed in cognitive psychology and theoretical linguistics still have their own integrity as involving non-derived and non-trivial semantic evaluability.

References

- Allport, A. 1988: What concept of consciousness? In A.J. Marcel and E. Bisiach (eds), *Consciousness in Contemporary Science*. Oxford: Oxford University Press, 159–82.
- Block, 1990: Consciousness and accessibility. *Behavioral and Brain Sciences*, 13, 596–8.
- Block, N. 1991: Evidence against epiphenomenalism. *Behavioral and Brain Sciences*, 14, 670–2.
- Block, N. 1992: Begging the question against phenomenal consciousness. *Behavioral and Brain Sciences*, 15, 205–6.
- Block, N. 1993: Review of D.C. Dennett, *Consciousness Explained*. *Journal of Philosophy*, 90, 181–93.
- Block, N. 1995: On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–87.
- Cannon, S.C. and Robinson, D.A. 1985: An improved neural-network model for the neural integrator of the oculomotor system: More realistic neuron behavior. *Biological Cybernetics*, 53, 93–108.
- Cannon, S.C., Robinson, D.A. and Shamma, S. 1983: A proposed neural network for the integrator of the oculomotor system. *Biological Cybernetics*, 49, 127–36.
- Chomsky, N. 1965: *Aspects of the Theory of Syntax*. Cambridge, MA.: MIT Press.
- Chomsky, N. 1976: *Reflections on Language*. London: Fontana/Collins.
- Chomsky, N. 1986: *Knowledge of Language: Its Nature, Origin, and Use*. New York: Praeger.
- Chomsky, N. 1990: Accessibility ‘in principle’. *Behavioral and Brain Sciences*, 13, 600–1.
- Churchland, P.S. and Sejnowski, T.J. 1992: *The Computational Brain*. Cambridge, MA: MIT Press.
- Clark, A. 1990: Aspects and algorithms. *Behavioral and Brain Sciences*, 13, 601–2.
- Davies, M. 1986: Tacit knowledge, and the structure of thought and language. In C. Travis (ed.), *Meaning and Interpretation*. Oxford: Basil Blackwell, 127–58.
- Davies, M. 1987: Tacit knowledge and semantic theory: Can a five per cent difference matter? *Mind*, 96, 441–62.
- Davies, M. 1989a: Tacit knowledge and subdoxastic states. In A. George (ed.), *Reflections on Chomsky*. Oxford: Basil Blackwell, 131–52.
- Davies, M. 1989b: Connectionism, modularity, and tacit knowledge. *British Journal for the Philosophy of Science*, 40, 541–55.
- Davies, M. 1990: Thinking persons and cognitive science. *AI and Society*, 4, 39–50. Reprinted in A. Clark and R. Lutz (eds), *Connectionism in Context*. London: Springer-Verlag.
- Davies, M. 1995: Two notions of implicit rules. In J.E. Tomberlin (ed.), *Philosophical Perspectives 9: AI, Connectionism, and Philosophical Psychology*. Atascadero, CA: Ridgeview Publishing Company, 153–83.
- Davies, M. and Humphreys, G.W. 1993: Editors’ Introduction. In *Consciousness: Psychological and Philosophical Essays*. Oxford: Blackwell Publishers, 1–39.

- Dretske, F. 1986: Misrepresentation. In R. Bogdan (ed.), *Belief: Form, Content and Function*. Oxford: Oxford University Press, 17–36. Reprinted in W.G. Lycan (ed.), *Mind and Cognition: A Reader*. Oxford: Basil Blackwell, 1990.
- Evans, G. 1980: Things without the mind. In Z. van Straaten (ed.), *Philosophical Subjects*. Oxford: Oxford University Press, 76–116. Reprinted in *Collected Papers*. Oxford: Oxford University Press, 1985.
- Evans, G. 1981: Semantic theory and tacit knowledge. In S. Holtzman and C. Leich (eds), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul, 118–137. Reprinted in *Collected Papers*. Oxford: Oxford University Press, 1985.
- Evans, G. 1982: *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, J. 1983: *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Grice, H.P. 1959: Meaning. *Philosophical Review*, 66, 377–88. Reprinted in *Studies in the Way of Words*. Cambridge, MA: Harvard University Press, 1989.
- Grice, H.P. 1989: *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Higginbotham, J. 1990: Searle's vision of psychology. *Behavioral and Brain Sciences*, 13, 608–10.
- Jackson, F. 1982: Epiphenomenal qualia. *American Philosophical Quarterly*, 32, 127–36. Reprinted in W.G. Lycan (ed.), *Mind and Cognition: A Reader*. Oxford: Basil Blackwell, 1990.
- Jackson, F. 1986: What Mary didn't know. *Journal of Philosophy*, 83, 291–95. Reprinted in D.M. Rosenthal (ed.), *The Nature of Mind*. Oxford: Oxford University Press, 1991.
- Levine, J. 1983: Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly*, 64, 354–61.
- Levine, J. 1993: On leaving out what it's like. In M. Davies and G.W. Humphreys (eds), *Consciousness: Psychological and Philosophical Essays*. Oxford: Blackwell Publishers, 121–36.
- Lloyd, D. 1990: Loose connections: Four problems in Searle's argument for the 'connection principle'. *Behavioral and Brain Sciences*, 13, 615–6.
- McGinn, C. 1988: Consciousness and content. *Proceedings of the British Academy*, 74, 219–39. Reprinted in *The Problem of Consciousness*. Oxford: Basil Blackwell, 1991.
- McGinn, C. 1989: Can we solve the mind-body problem? *Mind*, 98, 349–66. Reprinted in *The Problem of Consciousness*. Oxford: Basil Blackwell, 1991.
- Miles, F.A. and Lisberger, S.G. 1981: Plasticity in the vestibulo-ocular reflex: A new hypothesis. *Annual Review of Neuroscience*, 4, 273–99.
- Nagel, T. 1979: What is it like to be a bat? In *Mortal Questions*. Cambridge: Cambridge University Press, 165–80. Reprinted in D.M. Rosenthal (ed.), *The Nature of Mind*. Oxford: Oxford University Press, 1991.
- Nelkin, N. 1986: Pains and pain sensations. *Journal of Philosophy*, 83, 129–48.
- Nelkin, N. 1989a: Propositional attitudes and consciousness. *Philosophy and Phenomenological Research*, 49, 413–30.
- Nelkin, N. 1989b: Unconscious sensations. *Philosophical Psychology*, 2, 129–41.
- Peacocke, C. 1983: *Sense and Content*. Oxford: Oxford University Press.
- Peacocke, C. 1986: *Thoughts: An Essay on Content*. Oxford: Basil Blackwell.

- Peacocke, C. 1992: *A Study of Concepts*. Cambridge, MA: MIT Press.
- Perry, J. 1979. The problem of the essential indexical. *Noûs*, 13, 3–21. Reprinted in *The Problem of the Essential Indexical and Other Essays*. Oxford: Oxford University Press, 1993.
- Robinson, D.A. 1981: The use of control systems analysis in the neurophysiology of eye movements. *Annual Review of Neuroscience*, 4, 463–503.
- Robinson, D.A. 1987: The windfalls of technology in the oculomotor system. *Investigative Ophthalmology and Visual Science*, 28, 1912–24.
- Rosenthal, D.M. 1986: Two concepts of consciousness. *Philosophical Studies*, 49, 329–59.
- Rosenthal, D.M. 1990: On being accessible to consciousness. *Behavioral and Brain Sciences*, 13, 621–2.
- Rosenthal, D.M. 1993: Thinking that one thinks. In M. Davies and G.W. Humphreys (eds), *Consciousness: Psychological and Philosophical Essays*. Oxford: Blackwell Publishers, 197–223.
- Searle, J.R. 1983: *Intentionality*. Cambridge: Cambridge University Press.
- Searle, J.R. 1989: Consciousness, unconsciousness, and intentionality. *Philosophical Topics*, 17, 193–209.
- Searle, J.R. 1990a: Consciousness, explanatory inversion, and cognitive science. *Behavioral and Brain Sciences*, 13, 585–96.
- Searle, J.R. 1990b: Who is computing with the brain? *Behavioral and Brain Sciences*, 13, 632–40.
- Stich, S.P. 1978: Beliefs and subdoxastic states. *Philosophy of Science*, 45, 499–518.
- Strawson, P.F. 1959: *Individuals*. London: Methuen.
- Wilkes, K.V. 1988: ———, yìshì, duh, um, and consciousness. In A.J. Marcel and E. Bisiach (eds). *Consciousness in Contemporary Science*. Oxford: Oxford University