

Consciousness: Philosophical Issues

Volume 1, edited by E. Villanueva

[Matthew Elton](#)

Department of Philosophy
University of Stirling
Stirling FK9 4LA
UK

matthew.elton@stir.ac.uk

Copyright (c) Matthew Elton 1994

Received: July 1, 1994; Accepted: November 3, 1994

PSYCHE, 1(16), January 1995

<http://psyche.cs.monash.edu.au/v1/psyche-1-16-elton.html>

Keywords: consciousness, philosophy of consciousness, reductionism, qualia.

Review of: Enrique Villanueva (Ed.) (1991): *Consciousness: Philosophical Issues, 1*. Atascadero, California: Ridgeview Publishing Co. 276 p. Price: US \$25 pbk. ISBN: 0-924922-01-x.

1. Outline

1.1 The twelve essays in this collection are diverse, and, by the editor's own admission, five of them are not about consciousness at all. Nonetheless, taken together, and with one notable exception, they illustrate an emerging consensus on how to approach consciousness from a philosophical perspective. In this review, my general strategy will be to interleave an exposition of what I shall call the *reductive account* of consciousness with a commentary on the papers in the volume. Importantly, the reductive account moves away from the pessimism promoted by such authors as Nagel (1974), Jackson (1982), and McGinn (1991).

1.2 I shall begin my discussion with the papers by David Rosenthal and Daniel Dennett. These papers set out the key tenets of the *reductive account*, viz. that we can get a grip of intentional states of mind that is independent of consciousness, and that we can understand consciousness as a structure built out of such intentional states. Georges Rey and John Biro have rather different agendas, but their best arguments contribute to the *reductive* project. On the other hand, John Searle's paper is set fiercely against it.

1.3 Brian O'Shaughnessy's important contribution deserves special treatment. It is not clear quite where he stands with respect to reduction, but he offers an intriguing analysis of the structure of consciousness, of what the reductionist would seek to reduce. I shall close with some general comments on the volume and how it contributes to the evolving literature on the philosophy of consciousness.

2. Reductive Approaches

2.1 Rosenthal has been promoting his higher-order thought account of consciousness for some years now (e.g. 1986, 1993), and the paper in this volume is part of his on-going campaign. One of his main agenda items is that consciousness is not the mark of the mental. Rather, consciousness is what you get when mental states exhibit a certain structure. His account then introduces consciousness as "an extrinsic property of mental states" (p. 30). The story goes like this: a mental state, *p*, is conscious if the subject has an (occurrent) thought (that is based neither on observation nor on inference) to the effect that one is in mental state *p*. The presence of the higher-order thought makes the first-order mental state conscious. Of course, in the normal run of things, the higher-order thought is not conscious, but it can be if there is an occurrent thought directed upon it.

2.2 Now, some might regard Rosenthal's theory as providing a good account of self-consciousness or of introspection, but deny that it could account for phenomenal consciousness, for our everyday perceptual experience that is so full of life, character, and supposedly intrinsic 'raw feels.' So in this paper, 'Consciousness and Sensory Quality,' Rosenthal directly addresses the question of phenomenal consciousness. His aim is to show that states that are *about* sensory qualities, states representing red or pain, etc., need not be intrinsically conscious. We can, he claims, make good sense of the mind entertaining qualities of red or of pain quite independently of our being conscious of such qualities. And, of course, what it is for us to be conscious of such sensory qualities is to have a thought that we are in a mental state with the given sensory quality.

2.3 Rosenthal's argument turns on revealing an equivocation in the use of the expression "sensory quality" (or any of its many cognates: qualia, raw feel, phenomenal character). The expression is taken to mean both (i) what it is like to have a sensory state and also, on other occasions, (ii) the discriminatory role of the sensory state. If you run these two senses together, then, suggests Rosenthal, you will find it hard to accept his claim that there can be unexperienced sensory qualities. But for Rosenthal, it is quite proper to characterize the *quality* of mental states as a quality of pain or of red, even under reading (ii), the discriminatory role of the sensory state. Reading (i) only comes into play when we have a higher-order thought about being in the given sensory state. That the state has this quality, he urges, does not require that this quality be brought to the attention of consciousness. This seems plausible. We do speak of having the same headache all afternoon, even though the awareness of our pain is intermittent. And Rosenthal provides a good model for this. Throughout the afternoon the sensory state with the 'headaching' quality endures, but sometimes we have a higher-order thought directed at that state and sometimes we do not.

2.4 In his positive commentary on Rosenthal's paper, Dennett drives the point home when he remarks that "if there can be unfelt pains, there can be unconscious sensations of red" (p. 38). It is hard to find a natural example of unconscious colour perception, so Dennett suggests an artificial one. A video-game player may learn to associate a certain flashing red spot with danger. Now we can imagine that the player's attention is distracted in such a way that if the red spot flashes the player is unable to report this fact. Nonetheless, we may be able to gather evidence that 'red spot' information has been acquired on the basis of a galvanic skin response (which provides a measure of anxiety). I'm sympathetic to this approach. In fact, I'm sometimes surprised at how hard both Rosenthal and Dennett feel they need to work to establish the point, viz. that we can make sense of sensory quality independently of conscious apprehension of such qualities (though this may place me in the minority). However, this leaves a couple of areas of concern.

2.5 First, Rosenthal does not say enough about what having the higher-order thought does for the subject. If an unconscious pain can make you turn over in your sleep, why have conscious pains at all? If the unconscious state can influence behaviour, what is added by making it conscious? Looking at the neuropathology of blindsight (Weiskrantz, 1986) can be helpful here. Blindsight subjects fail to have visual experiences in a part of their visual field. But some information about the contents of this part of the visual field is clearly getting through, since if you ask a patient to make guesses about the field's contents they do much better than chance. Now, Rosenthal can explain this by suggesting that the sensory state concerned with the 'blind' part of the field is present, but, because of damage to the brain, that sensory state cannot engender a higher-order thought. What this seems to mean is that the role such sensory states can play in behaviour is severely curtailed. Blindsight patients cannot spontaneously respond to stimuli in the blind region -- they need to be prompted to guess. They cannot recruit information about that region in the planning of actions, and they can only come to learn of the sensory state by indirect means. Clearly there is plenty of scope for Rosenthal to extend his analysis of higher-order thoughts and to show how they relate to certain kinds of mental activity. In particular, higher-order thoughts seem to be central to a capacity to report sensory qualities, to exploit them in making plans, to responding spontaneously to (non-urgent) stimuli, and perhaps even to constructing one's own self-narrative or mental biography.

2.6 Rosenthal does begin to address some of these issues towards the end of his paper, where he discusses the relationship between higher-order thoughts and personhood. (See, Frankfurt, 1971, for a good discussion of this relationship.) He also claims that "[relatively] weak conceptual resources will suffice for a higher-order thought to refer to one's own sensory states" (p. 32). This allows Rosenthal to grant sensations to animals and neonates. But his generosity has limits as he claims that "we have no reason to suppose that animals other than persons [1>](#) are aware of whatever higher-order thoughts they may have" (p. 35). And, for Rosenthal, it is these 'higher-higher-order' thoughts along with "some measure of rational connectedness" that are crucial to personhood. These kinds of issues are critical but, sadly, Rosenthal displays only his intuitions and gives us no firm argument to back them up.

2.7 My second worry I raise on behalf of the traditional pessimist, who may claim Rosenthal has just left something out of his account. He has only discussed the psychological aspects of consciousness, only talked about what Block (in print) calls "access" consciousness and not talked about what really matters, "phenomenal" consciousness, the home of qualia and raw feels. Surely phenomenal consciousness does not simply amount to an occurrent thought that one is in the kind of mental state typically caused by smelling the scent of a rose? Surely the content of the mental states is one thing and the quality another. Content alone cannot suffice to account for experience. Rosenthal, and his co-conspirator Dennett, are clear on this point. The position is spelled out by the latter in the following bold statement:

Don't our internal discriminative states also have some special "intrinsic" properties, the subjective, private, ineffable properties that constitute the *way things look to us* (sound to us, smell to us, etc.)? No. The dispositional properties of those discriminative states already suffice to explain all the effects: the effects on peripheral behaviour (saying "Red!", stepping the brake, etc.) and "internal" behaviour (judging "Red!", seeing something as red, reacting with uneasiness or displeasure if, say, red things upset one). Any additional "qualitative" properties or qualia would thus have no positive role to play in any explanations, nor are they vouchsafed to us "directly" in our intuition. Qualitative properties that are intrinsically conscious are a myth, an artifact of misguided theorizing, not anything given pretheoretically (p. 40).

2.8 I'm eager to agree with these sentiments. As it happens, I do think that episodes of consciousness can be explained by reference to episodes of thought, thoughts which have content, but which do not have some additional character or intrinsic quality. Experiences of red, for me, just are judgements that there is something red out there along with something like a judgement that I've made just such a judgement (that there is something red out there). So I subscribe to the broad class of views that attempt to *reduce* claims involving consciousness to claims involving content. The worry here and, to a lesser extent, elsewhere (Dennett, 1991) is finding a positive argument for this reductionist view.

3. Arguing for Reduction

3.1 Actually, it is hardly surprising that there is no positive argument for reductionism. The arguments of the pessimists, if they show anything, show that such is not available. Recall the 'knowledge argument' developed by Jackson (1982). If Mary were an expert on hydrogen, oxygen, and chemistry, then she would be able to find out all about the macroscopic properties of water. But, from the point of view of explanation, consciousness is not like that. Give Mary all the facts there are to know about neurophysiology and psychology (including the content of judgements entertained by subjects) and Mary, merely in virtue of having these facts available, won't come to have new experiences. And, because of this logical point (but see Lewis 1990), there is some kind of explanatory gap in the case of consciousness. (This interpretation of the

knowledge argument is rehearsed by Rey (p. 102) *and* Biro (p. 128).) What the reductionist tries to show is that there is no gap in ontology. Just because a certain set of facts does not obviously entail consciousness, does not mean that those very facts do not fix consciousness, does not mean that they don't explain it as best it can be explained.

3.2 So the reductionist strategy goes something like this:

- (1) Provide an account of the mental that explains the phenomenology of experience or, at the very least, explains verbal reports and behaviour.
- (2) Show how this account predicts the kinds of concerns raised by the pessimists, e.g. the explanatory gap.
- (3) Show how alternative accounts that call for extra ingredients are false or, at least, very unattractive. (By extra ingredients I mean something over and above mental states with content, e.g. mental states with intrinsic character or biological properties of a cognitive mechanism.)
- (4) Claim that, after all this work, the reductive account wins by default.

This is not the most forceful of argumentative strategies, but, in the case of consciousness, it looks as though it is the best you can get.

3.3 In their contributions to this volume, Rosenthal and Dennett go some way to addressing the first and third points, but say very little about the second. Meeting the first two points involves meeting what can be dubbed the 'polemical challenge':

A polemically successful answer to the [problem of consciousness] and the knowledge argument must explain the intuitive appeal of the contention that what it's like is irreducibly non-physical information about experience (Nemirov, 1990, p. 495).

And for "non-physical" we can read "non-psychological" as well, since the knowledge argument clearly stands against what I am calling the reductive account. Fortunately for the reductionist, however, the papers by Biro and Rey forcefully address the polemical challenge. (Rey is a reductionist, I think. This is less clear in the case of Biro.)

3.4 Let's take a look at these two papers. Rey's "Sensations in a Language of Thought" comes in two parts. First Rey defends a modified version of the language of thought hypothesis (Fodor, 1987) which is carefully tweaked to cope with perceptual experience as well as linguiform thinking. This account is Rey's take on naturalizing the intentionality of perceptual judgements. I've nothing to say about that here, although I do have grave doubts about such approaches. The second part deals with various aspects of the polemical challenge. Rey shows how a special class of inner sentences (specialized for dealing with perceptual information) can account for phenomenology and will give rise to "privacy", "unanalyzability", "lack of grain", and "ineffability." For example,

privacy is accounted for by the fact that a subject's own sensory systems are intimately involved in the production of the relevant special inner sentences (pp. 97-98). And "the ineffability" of sensory experience is a consequence of the fact that no expressions in a natural language come close to playing the specific role that the [special sensory-sentences] play in a system's internal language" (p. 99).

3.5 Rey's arguments strike me as effective, if not entirely novel (see, for example, the bracing discussion of these issues in Dennett, 1969, chapter 7, and 1978); but I do find it strange that he feels the need to link them to his language of thought based account of content. The two issues seem to be largely orthogonal. The arguments he makes for special sentences is equally applicable to sub-personal content (see below) that is not identified with representations in a computational account of cognitive function. What can be said, however, is that he shows quite clearly that what is not wrong with language of thought based accounts is that they cannot account for consciousness.

3.6 Biro's paper, "Consciousness and Subjectivity", aims to address the following complaint:

Consciousness is essentially theory-resistant: this is because its essence is bound up with *subjectivity*, something that by its very nature must elude theorizing, since the latter's essential aim is to give an objective account of its phenomena (p. 115).

3.7 Biro notes that the expression 'point of view' often seems to do most of the work in arguing for the irreducibility of the subjective; especially in Nagel's work. What does the expression 'point of view' mean? Biro offers us three suggestions. First, it can refer to "the beliefs, conceptual framework or even values of some subject or group" (p. 117). Second, it can be taken to denote a location or vantage point, as in the view from the top of the tower. Biro calls this the 'fixed' reading and contrasts it with a third suggestion, the 'portable' reading. This reading indexically refers to the point of view of the subject, wherever the subject happens to be right now. It is portable since it moves as the subject moves.

3.8 The first and second reading clearly present no special problems for objective theorizing. The third does present a problem but, argues Biro, it is a trivial one:

If what matters about my experience is its mine-ness..., its being so does indeed seem to be the sort of things that cannot be included in any description of the experience, and thus it may really make experience theory resistant. The trouble is that it does so in a merely trivial way... Mere ownership does not in general contribute anything to the nature or character of the thing owned, and there is no reason to think that it does so when that thing is experience (pp. 121-122).

3.9 Biro neatly shows how Nagel slides between the portable and fixed readings, and argues that once the distinction is made clear Nagel's dramatic conclusions fall. (He also

makes some interesting remarks about types of point of view, e.g. about the differences between bat and human points of view. Here I think Biro underestimates the importance of the difficulty in drawing a line between a cognitive and perceptual system -- this difficulty, I would suggest, is conceptual rather than empirical.) The overall conclusion is that although different subjects do have different points of view on the world, this does not license pessimism with respect to objective theorizing. An individual's point of view is highly specific (as informed by background beliefs, type of cognitive and perceptual system, etc.) and rich with indexically specified content, but it does not comprise any "inexpressible additional content ... that is in principle unknowable by another individual" (pp. 130-131).

4. Against Reduction

4.1 The reductive account turns on making good sense of the attribution of intentional states without first invoking the concept of consciousness. Searle thinks this is just daft:

The attribution of any intentional phenomena to a system, whether "computational" or otherwise, is dependent on a prior acceptance of our ordinary notion of the mind, the conscious "phenomenological" mind (p. 47).

4.2 Searle is anxious to qualify this strong claim, since he doesn't want to deny unconscious mental states, which, for him, include repressed beliefs and desires as well as intentional states that are simply not present to the mind at a particular time. He argues that unconscious mental states are best understood as dispositions of the underlying neurophysiology of the brain to give rise to occurrent conscious mental states. So consciousness, even here, remains the mark of the mental.

4.3 The argument Searle makes turns on two features of intentional states, two features that he claims require us to invoke consciousness.

[Intentional] states, conscious or unconscious, essentially have aspectual character and this aspectual character is essentially subjective, in the sense that it cannot be exhaustively accounted for in third person "objective" terms (p. 58).

The aspectual character of an intentional state is, very roughly, its sense or cognitive significance. When I'm thirsty my thought is directed at water (my thought *refers* to that stuff). But I think of water under the aspect 'water' (the liquid stuff that is all around) rather than, say, H₂O (a chemical described in a serious textbook). I wholly agree that aspectual shape is a mark of an intentional state, and I also agree that many thoughts have an indexical character (e.g. because they are about what *I* am looking at *here* and *now*), which gives rise to a certain kind of subjectivity. But I am wholly unpersuaded by his arguments that this subjectivity "cannot be exhaustively accounted for in third person "objective" terms" (ibid.). Here Searle seems to be equating a puzzle about phenomenal experience (qualia) with a puzzle about intentionality. Yes, it is clear that he thinks the

two are inextricably linked. What is not clear is how he supports this claim with argument.

4.4 I don't think bacteria or insects are conscious, and I don't think Searle does either. But the best way to explain the behaviour (rather than the mere bodily movements, see Dretske 1988) of such creatures is to ascribe cognitive and motivational states that are *about* features of their world. These states will be subjective (not in Searle's inflated sense, but in the defensible indexical manner) and will have aspectual shape. For example, a simple animal might exploit a certain chemical trace as a reliable indicator of food nearby. Such a creature directs its desires (motivational states) at its food source (this is what the motivational state *refers* to), but it will ignore the food if some devious experimenter has suppressed the chemical trace. Its desire comes to be *about* its food source only via its intentional state and the way in which it can represent the world.

4.5 Although Searle makes much of the conscious-unconscious distinction, he is not prepared to countenance a distinction between sub-personal and personal content (Dennett, 1969). Searle's conscious and unconscious states are both examples of personal content, content for the unified and conscious agent (a notion which Searle, unlike Rosenthal and O'Shaughnessy, does precious little to analyze). Sub-personal content, on the other hand, does not require us to invoke a unified and enduring psychological subject. We can use sub-personal content to explain the behaviour of simple creatures^{2>} and to explain parts of the behaviour of our own cognitive systems. Think about the judgements that your visual system makes when it plans and executes the movements of your eyes as they saccade across a page of text. If we want to explain the reasons for, rather than the causes of, such movements, we have to invoke intentional language--we have to invoke content. But such content is not available to the person, and is probably not even available in principle. Such content is sensibly called sub-personal. If one can describe such sub-personal intentional states as thoughts, then they certainly aren't 'thoughts' of the subject whose eyes are being directed. (They are more like 'thoughts' of the subject's brain.)

4.6 Once a distinction is made between personal and sub-personal content, it is possible to embrace many of Searle's points by taking them to apply only to personal content. The reductionist strategy is to show how sub-personal contents conspire to create personal content by bringing into being thoughts that are about an on-going psychological entity with memories, intentions, and phenomenal sensations.

5. What Is Being Reduced

5.1 The star paper in the collection is O'Shaughnessy's "The Anatomy of Consciousness". I can hardly do it justice in the short space available here, but I shall try to give you its flavour. Unlike the other essays I have discussed, this paper is not directly concerned with giving an account of intentionality or phenomenal experience. Rather:

It is [concerned] with the vastly familiar light that appears in the head when a person surfaces from sleep or anaesthetic or dream. In other words

with the state we call "waking", which I shall mostly refer to as "consciousness" (p. 135).

5.2 The paper begins with an investigation into the logical form of states of consciousness. This then broadens into a discussion of the properties of the various modes and the manner in which they interrelate. O'Shaughnessy argues that consciousness (being awake and alert) is the essential defining mode of all other modes of consciousness, such as sleep and dreaming, and that other modes are "privative derivatives". Consciousness, for a particular creature, is the mode of being such that, given the capacities of the creature, it maximises the degree to which the creature is in touch with reality. For example, it is the mode of being in which "belief-formation out of veridical perception should be such as normally to make *knowledge* of that belief" (p. 157). This, O'Shaughnessy argues, is the "prime function of the state of consciousness" (p. 159).

5.3 Of course there is something it is like to be a conscious creature because consciousness supports episodes of thought and phenomenal experience. But O'Shaughnessy insists that it makes no sense to claim that consciousness itself is like anything, or that it is directed upon anything in an intentional manner. And so, curiously, consciousness has a kind of invisibility, on account of the fact that "consciousness *is* a system and *supports* a system" (p. 169). And this makes investigating its structure somewhat difficult. However, investigating the privative modes of consciousness can help. The various privative modes exhibit absences of essential component functions of consciousness and so help make plain those component functions that are otherwise difficult (though not, if I read O'Shaughnessy aright, impossible) to divine.

5.4 The various privative modes are, like consciousness, stable systems for underpinning psychological activity, but modes which fail to fulfil the prime function of the state of consciousness in that they fail to aim squarely at reality. I can best illustrate this idea by discussing one of O'Shaughnessy's key examples: dreaming. He uses this example to explore both the privative mode, dreaming, and also to reveal the role of action and temporal thought in consciousness (pp. 161-167). He argues that in the dream the subject's thoughts of the future are directed not at the future of the dream, but at the future of reality. In our nightmares we fear genuine danger, not a dream-like danger. And with respect to the past:

The dream present lacks modes of solidarity with, or responsibility to its past... I can be dreaming of *anything* in one instant and *anything* in the next. In a sense therefore the dream is created anew in each instant: not to the point where narrative disintegrates into unsynthesizable fragments; but in so far as the character of the present experience necessitates neither a past nor a cognitive attitude to the past... [It] is continuity of content, rather than persistence of contents, that unifies the dream (p. 165).

5.5 In sum "the dream present is a sort of Time Island" (p. 166) that lacks proper connection with a past and a future. The study of the privative mode of dreaming helps us understand the unprivated state. For:

[The dream] relation to Time is inconsistent with consciousness. This is because consciousness requires that we be in a position to perceive *events across time*; which is to say, not merely events which themselves cross time, but the very profile across time of the event itself (p. 166).

5.6 O'Shaughnessy develops this point by asking us to consider what goes on when we understand a spoken sentence, or listen to a piece of music. "[Such] perception is only possible if we retain cognitive links with our internal past, and an open but real cognitive connection with an anticipated internal future." With the aid of some elaborate, and possibly quite delicate, argument O'Shaughnessy suggests a link between internal mental action (the will) and temporal capacity:

[We] will not achieve [consciousness of the spatio-temporal world] if we do not relate cognitively and experientially to the past and future of things; and we will not relate thus to their past and future if we do not relate thus to our own internal past and future; and that this last is something that is manifestly accomplished in intentional internal action... (p. 167)

5.7 So the picture of consciousness painted is one of an active and enquiring mind that aims to keep itself properly in touch with reality. The various component parts of consciousness are deeply interrelated, but we can come to understand them through careful analysis and by considering various privative modes. It is this structure, then, that the reductive account seeks to explain. It seeks to show how sub-personal contents can conspire to give rise to the system of consciousness, and its derivative modes, that O'Shaughnessy describes. It is not clear to me what he would make of such a project. He says of the system that is consciousness that:

[Here] mental explanation comes to an end, and openly gives ways to the cerebral. While *the properties* of consciousness are susceptible of mental explanation in terms of its own mental constituents, the *existence* of consciousness is not; and the same must be true of the constituents themselves. Consciousness and constituents arise simultaneously together poised on the one purely physical base (pp. 169-170).

5.8 Villanueva (introduction, p. 7) interprets this remark as opening up space for a "sub-personal, cerebral" approach. The supporter of the reductive account is unlikely to be happy with the implied interchangeability between the expression 'sub-personal' and the expressions 'cerebral' or 'purely physical base.' Whether sub-personal content can be reduced to the physical (or the computational) is a further question, and one which is independent of the reductive account's reductive aspirations. What is not clear to me in O'Shaughnessy is whether he thinks there is no intermediate level between consciousness (personal content) and the physical, or whether he thinks that the intermediate level (sub-

personal content) cannot itself be reduced. But, O'Shaughnessy's views aside, the supporters of the reductive account can exploit his excellent analysis for their own ends. And this analysis, in part, fills in some of the gaps in Rosenthal's account of which I complained above.

6. Closing Remarks

6.1 This volume comprises some useful essays on consciousness, as well as five essays on other philosophy of mind topics, which I've not discussed here.^{3>} I've concentrated on certain themes, but there's much more material to be mined from this collection. On the other hand, most of the content here can be found in other places (a good source is Davies & Humphreys, 1993), with the exception, to my knowledge, of the excellent O'Shaughnessy paper. Added to this the very poor standard of editorial production (an unacceptable number of typographical errors, missing references, inaccurate cross-referencing, etc.), this is not a volume I could whole-heartedly recommend.

6.2 For some time much of the philosophical literature on consciousness has focused on debating the very possibility of providing an account, with much of the lead being taken by various prominent pessimists. The papers in this volume indicate a new way for the debate to turn. Optimists, especially Dennett and Rosenthal, have something like a positive account of consciousness to offer, the reductive account. The very coherence of such an account can be questioned, as it is by Searle, and the details can be hotly disputed. And, of course, the reductive account almost certainly won't cheer the diehard pessimist, the anti-physicalist, or the epiphenomenalist, all of whom think that what makes conscious states light up is something outside the realm of scientific analysis. Their pessimism won't just go away, but there is still something in the new turn of debate for them. For although adherents of the reductive account aim to provide necessary and sufficient conditions for consciousness, in the interests of keeping the discussion fresh, and avoiding the tiresome gloom of the pessimistic school, reductionists may choose to couch their theory as one of merely necessary conditions. Certainly all parties can be interested in what it is the reductionist seeks to reduce. O'Shaughnessy explores this territory in an unashamedly *a priori* fashion (p. 170), but more empirical approaches could also be offered, and the merits of the contrasting methodologies debated. These kinds of issues, issues that are put into focus by putting forward a positive account of consciousness, take the philosophy of consciousness into more interesting areas than the largely negative debate fostered by the important work of the pessimists.

Acknowledgements

This work was made possible by SERC grant ref. 91313730.

Notes

1. Rosenthal ought to use the word 'human' here, since the use of the word 'person' in contrast with animal simply begs the question he is addressing.

2. I would claim that such organisms do not realise psychological subjects at all -- such organisms bring forth biological selves, but not psychological ones.

3. For the record the remaining five begin with an indifferent paper on internalism and externalism by Ernest Sosa. Then, there is an interesting paper by Donald Davidson, titled "What is Present to the Mind." Davidson discusses the problems faced by an externalist theory of self-knowledge. Akeel Bilgrami provides a good commentary on Davidson. Finally, James Tomberlin has a paper that, in a less than edifying manner, discusses Chisholm's views on intentionality and self-ascription. Lynn Pasquarella's excellent commentary on Tomberlin describes Chisholm's and Tomberlin's views in a clear and concise way--she also demonstrates how Chisholm's account readily withstands Tomberlin's criticism.

References

Block, N. (in print). On a confusion about a function of consciousness. *Behavioural and Brain Sciences*.

Davies, M. & Humphreys, G. (Eds.) (1993). *Consciousness*. Oxford: Blackwell.

Dennett, D. (1969). *Content and consciousness*. London: Routledge & Kegan Paul.

Dennett, D. (1978). Two approaches to mental images. In *Brainstorms* (pp. 174-189). Hove: Harvester.

Dennett, D. (1991). *Consciousness explained*. Boston: Little, Brown and Company.

Dretske, F. (1988). *Explaining behaviour: reasons in a world of causes*. Cambridge, MA: MIT Press.

Fodor, J. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.

Frankfurt, H. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy*, 68, 5-20.

Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127-136.

Lewis, D. (1990). What experience teaches. In W. G. Lycan (Ed.), *Mind and cognition: a reader* (pp. 499-519). Oxford: Blackwell.

McGinn, C. (1991). *The problem of consciousness*. Oxford: Blackwell.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review* 83, 435-450.

Nemirov, L. (1990). Physicalism and the cognitive role of acquaintance. In W. G. Lycan (Ed.), *Mind and cognition: a reader* (pp. 490-499). Oxford: Blackwell

Rosenthal, D. M. (1986). Two concepts of consciousness. *Philosophical Studies*, 49, 329-359.

Rosenthal, D. M. (1993). Thinking that one thinks. In M. Davies and G. W. Humphreys (Eds.), *Consciousness* (pp. 197-223). Oxford: Blackwell.

Weiskrantz, L. (1986). *Blindsight*. Oxford: Clarendon.