Christopher Mole

Attention to Unseen Objects

Abstract: *Can one pay attention to objects without being conscious of them? Some years ago there was evidence that had been taken to show that the answer is 'yes'. That evidence was inconclusive, but there is recent work that makes the case more compellingly: it now seems that it is indeed possible to pay attention to objects of which one is not conscious. This is bad news for theories in which the connection between attention and consciousness is taken to be an essential one. It is good news for theories (including mine) in which the connection between attention and agency is taken to be essential.*

1. A Point of Contention

According to one view of consciousness, and of its relationship to attention, a conscious person will normally be conscious of several different things. Some proper subset of these things can be phenomenally foregrounded and cognitively prioritized. The things in that subset will, in virtue of this foregrounding and prioritization, be the things to which that person is paying attention. Such a view entails that unconscious attention is an impossibility, on a par with the impossibility of being the star without appearing in the show. Attention, according to this view, *just is* the foregrounding and prioritization of consciousness.

In a 2008 article I suggested (1) that this view of the relationship between attention and consciousness is quite a natural one, and (2) that the evidence that had been taken to refute it did not do so (Mole, 2008a). Both claims have come under attack in the subsequent literature: evidence has been adduced to show that this view is not as

Correspondence:
Email: chris.mole@ubc.ca

Journal of Consciousness Studies, **21**, No. 11–12, 2014, pp. 41–56

natural as I had supposed (de Brigard, 2010), and to show that the view is not as defensible as I had suggested (Kentridge, de-Wit and Heywood, 2008; Koch and Tsuchiya, 2008). The present article is concerned with the second of these points. It explains that there are now reasons for giving up the claim that I once thought natural. It then considers the philosophical repercussions of doing so.

The claim that is at issue might be thought of as the claim that consciousness is necessary for attention. That would not be an inaccurate way to think of it, but — because of the ambiguity between 'creature consciousness' and 'state consciousness' (Rosenthal, 1986; White, 1964) — I do not think that it is sufficiently precise. The claim at issue is better formulated as follows:

> For all persons and all things, if the person is attending to the thing then the person is conscious of that thing.

This formulation was referred to, in my earlier treatment, as 'α'. It will be convenient to continue with that convention here.

If α were true then the things to which one is paying attention would be a subset of the things of which one is conscious. The door would then be open to the pursuit of a consciousness-first strategy, when one is giving an explanation of attention. One could start with consciousness — supposing it to be explained elsewhere — and could explain what attention is by specifying how some proper subset of the conscious field comes to be foregrounded, so that it qualifies as a locus of attention.

If α turned out to be false, on the other hand, then any such strategy for the explanation of attention would no longer be satisfactory, since it would fail to account for those instances of attention that occur unconsciously. But, because of that, the door would then be open to the pursuit of an attention-first strategy, when one is in the business of explaining consciousness.

No more than one of these explanatory strategies can be viable. Both have their advocates.

The consciousness-first strategy for explaining attention has been adopted by Declan Smithies, and by Sebastian Watzl. Smithies has suggested that the thing that needs to happen to consciousness, in order for it to constitute attention, is for the contents of consciousness to be available for 'rational access' (Smithies, 2011). Watzl has suggested that consciousness comes to be attention-constituting when a structure is imposed upon it in which some parts are, in a sense that he tries to make precise, more *peripheral* than others (Watzl, 2011). Both

treat attention as if it were a particular kind of consciousness. They therefore require α to be true.

The contrary strategy — i.e. the attention-first strategy for explaining consciousness — has been adopted by Jesse Prinz. It is pursued in his 2012 book, *The Conscious Brain* (Prinz, 2012), and in a series of associated articles (Prinz, 2010; 2011). According to Prinz's 'Attended Intermediate Representation' theory (which he abbreviates to 'AIR'), consciousness is what happens when attention is paid to stimuli that are represented in a certain 'intermediate' way. Prinz's project therefore depends on rejecting the claim that attention is just one particular variety of consciousness (since if it were then he would be using one particular form of consciousness to explain consciousness itself, and the resulting explanation would be circular). He must therefore reject the approach that Watzl and Smithies have taken.

My main business in the present article will be to review the most impressive of the experiments that now provide evidence for the falsity of α. Such evidence is, for reasons that have already been indicated, bad news for the explanatory project that has been pursued by Smithies and Watzl. We shall be seeing that it is also unwelcome news for Prinz's project, as Prinz himself understands it.

In so far as there is good news here it will be good news for a theory such as that advocated by Wayne Wu (2011) — or, indeed, by me (Mole, 2005; 2011). According to these theories attention is not essentially related to consciousness (although it might nonetheless be essentially related to psychological processes in which consciousness is normally, but contingently, involved). In these theories attention is understood as being essentially related, not to consciousness, but to *agency*, broadly construed. I conclude with some remarks on the significance of this.

2. Evidence Against α

If α were true then it would be impossible to attend to a thing without being conscious of that thing. The empirical case against α comes from evidence that this is not impossible, since there are actual cases of it.

The clearest of these cases seemed, back in 2008, to be the case of a blindsight patient, GY (described in Barbur, Ruddock and Waterfield, 1980). A series of experiments by Robert Kentridge, Charles Heywood, and their collaborators had found GY to be capable of selectively attending to locations that fall within the blind part of his visual field (Kentridge, Heywood and Weiskrantz, 2004; Kentridge, Nijboer

and Heywood, 2008; Kentridge and Heywood, 2001). That was an important finding to have made, but several philosophers and psychologists were too quick to suppose that it provided a clear counter-example to a claim like α. Although Kentridge and Heywood's experiments show that GY can differentially direct his attention to different parts of the space that falls within the blind part of his visual field, it is not clear that their experiments show GY to be directing his attention to the unseen *objects* that are located in those parts of space. GY is not conscious of the objects that are presented on his right hand side, but he may not be able to pay attention *to those objects*. It may simply be that, when he focuses his attention on the parts of space in which the objects happen to occur, his processing of information pertaining to those objects is facilitated. If this is the correct description of GY's attention to his blindfield then he is not a counter-example to α.

This alternative characterization of GY's attention depends on the drawing of a distinction between attending to an object and attending to the space within which that object occurs. The tenability of such a distinction is at the crux of the current discussion. The distinction is a particularly difficult one to operationalize in the case of the blindsighter. It is partly as a result of this difficulty that the empirical case against α has moved, in the more recent literature, away from the consideration of unconscious attention in blindsight, and towards the consideration of unconscious attention in the normal population.

The move from considering a blindsight patient to considering normal subjects brings with it the need to contrive special stimuli, of a sort that can be presented while subjects remain unconscious of them. Psychologists have a well-established repertoire of tricks for creating such stimuli, but the attempt to demonstrate attention in the absence of consciousness has led to a new trick being developed. This trick — which was first used in a 2013 paper by Liam Norman, Charles Heywood, and Robert Kentridge — depends on presenting one's experimental participants with shapes such as the rectangles in Figure 1, the edges of which are defined purely by contrasts of visual texture.

In the experiment that first used this trick, Norman, Heywood and Kentridge's (2013) participants were presented with a screen showing a grid that was filled with Gabor patches (except that sixteen evenly spaced cells of the grid were empty). The Gabor patches could be oriented vertically or horizontally, and this difference between the orientation of the patches was used to define two rectangles, as it does in Figure 1.

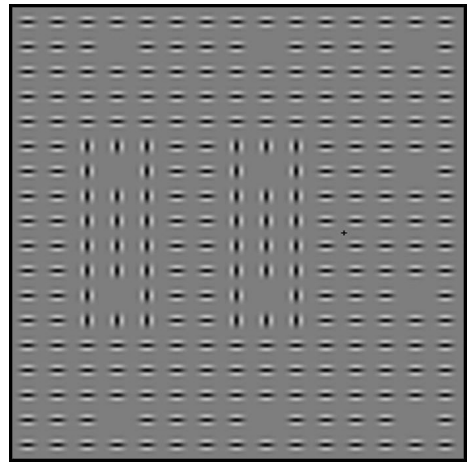

Figure 1. An example of one time slice of Norman, Heywood and Kentridge's stimuli. The two rectangles are defined purely by the orientation of the Gabor patches. Image kindly provided by Liam Norman.

The clever part of Norman, Kentridge and Heywood's trick is that, if all of the vertical patches switch to being horizontal, and all of the horizontal patches switch to being vertical — and if they do this switching back and forth at the rate of one switch every 30 ms, then the patches all become a blur. The result is that the rectangles become completely invisible: when subjects are given a signal-detection task, and are asked to make a guess about whether the rectangles are present or absent, their guesses contain no information about whether the rectangles are there or not.

Norman, Heywood and Kentridge's remarkable discovery — the details of which we shall be considering shortly — is that the participants in this experiment can be prompted to pay attention not only to the parts of the screen on which these invisible rectangles occur, but to the invisible rectangles themselves. It is this that shows, contrary to my own earlier suggestion (Mole, 2008a,b), that claim α must be mistaken. Norman, Heywood and Kentridge's (2013) experiment shows that it *is* possible to attend to a thing without being conscious of it.

In order to show that the participants in this experiment are directing their attention to the invisible rectangles themselves, and not merely to the space in which those rectangles occur, Norman, Heywood and Kentridge depend on a variation of the same-object advantage. It is worth reminding ourselves of what this is, since it is the same-object advantage that enables us to operationalize the distinction between

attention to locations and attention to the objects that fall within those locations.

It should be uncontroversial that there is some such distinction to be drawn. The claim that subject *S* is attending to region *R*, and the claim that object *O* falls within region *R*, do not together imply that *S* was attending to *O*. To show this, let *O* be something imperceptible, or something that behaves in such a way as to elude attention (such as the flowing stimuli used by van Marle and Scholl, in their 2003 paper). Alternatively (and as I have suggested previously), let *R* be the retinal blind spot, and *O* be a mark that falls wholly within it. In these cases there can be attention to the location within which an object falls, without the attentive subject thereby paying any attention to that object. There must therefore be *some* distinction to be drawn between attending to a location, and attending to an object in that location.

Some commentators have suggested that it is specious to take the preceding argument as establishing anything more than a *logical* distinction between attention to objects and attention to locations (Koch and Tsuchiya, 2008, p. 45). It must be admitted that our earlier discussion required us to draw this distinction in what might seem to be an especially problematic case. We wanted to distinguish between attending to the things in a location and attending to the location of those things even in GY's case, in which paying attention to a location facilitates the processing of information that pertains to the things in that location. Here one might be tempted to complain that the drawing of our distinction makes no difference. Even if there is a logical distinction between attending to a thing and attending to the part of space containing that thing, it might be said that — in those cases where attention is paid in such a way as to facilitate one's processing of information pertaining to the things in the space that has been attended — there is no *psychological* distinction between attending to a thing and attending to its part of space. It might then be said that the distinction between attention to locations and attention to the things in those locations is worth drawing only when attention to a region of space has no effect on the processing of information from the things that fall into that region, so that we only really need this distinction when we are considering exceptional regions, such as the retinal blind spot, or exceptional things, to which attention cannot normally be paid.

There is something right about this line of complaint, and we shall try to do justice to it in what follows, but in the form just given the complaint clearly goes too far. Psychologists do need to distinguish between attention to locations and attention to objects, even in the normal case. The 'same-object advantage' shows why. It shows one

way in which the attention to objects/attention to locations distinction can be drawn, even for unexceptional regions and unexceptional objects, and even when attention to a space does facilitate the processing of objects that fall within it.

The existence of the same-object advantage was established in a now-classic study, conducted by Robert Egly, Jon Driver and Robert Rafal (1994). Egly, Driver and Rafal presented their participants with a pair of rectangles like those shown in Figure 2. Unlike in the Norman, Kentridge and Heywood (2013) experiment, there was nothing special about these rectangles. They were defined by a light grey line on a dark grey ground and were presented, sometimes vertically and sometimes horizontally, for a period of seconds, on a screen that was in clear view. Participants were instructed to fix their gaze on a mark lying halfway between these rectangles, in the middle of the screen. The experimenters then changed the colour of one end of one of the rectangles for 100 ms, as an attention-directing cue. The participants had to respond as quickly as possible to a square target stimulus, which sometimes occurred and sometimes didn't occur, 200 ms after this attention-directing cue had been removed.

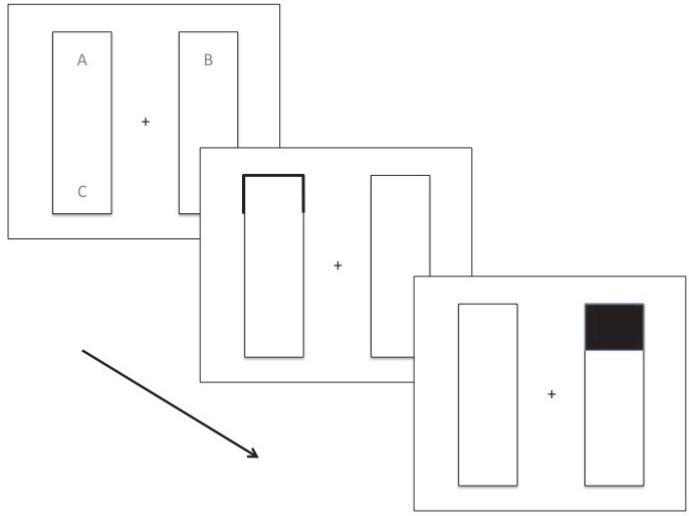

Figure 2. A schematic representation of one trial in Egly, Driver and Rafal's demonstration of the same object advantage. When the top left location is cued, reactions to targets in the top right are faster only if the bars are oriented horizontally, so that the cue and the target locations fall within the same object. The attention-grabbing cue therefore draws attention to the object of which it is a part, and not merely to the location in which it occurs.

In order to gauge the way in which the participants in this experiment were paying attention, Egly, Driver and Rafal (1994) timed their reactions to the target squares, on the assumption that the reaction times would be quicker if the cue had directed the participants' attention to the location in which the target then appears, and would be slower if that cue had directed their attention to a location other than the location of the target.

The distance between the rectangles was, crucially, the same as the distance from one end of a rectangle to the other: the location marked 'A', in Figure 2, is as far from the location marked 'B' as it is from the location marked 'C'. If the rectangles were oriented horizontally then the two locations that would previously have fallen into the top (or bottom) of two different rectangles now occurred at opposite ends of the same rectangle. These stimuli therefore allowed Egly, Driver and Rafal to distinguish between the effect of directing attention with a cue that occurs *near* to the subsequent target, and the effect of directing attention with a cue that falls *within the same visual object* as that target. By showing that a cue in one location draws attention to another location only if the two locations fall within the same visual object, their experiment reveals that a 100 ms cue does not simply pull attention to the part of space that that cue occupies. Such a cue pulls attention to the *object* of which it is a part. The same object advantage therefore illustrates the distinction that we are interested in. The presence of this advantage indicates that attention is being paid to an object, and not just to the space in which that object falls.

The point to notice here is that the same-object advantage gives sense to the distinction between attention to objects and attention to their locations, even in the case where the region in question is unexceptional, where the objects are normal candidates for attention, and where they are being processed in an attention-enhanced way. It is not simply that the rectangles in Egly, Driver and Rafal's experiment fall within the attended area, nor is it simply that, in virtue of falling into that area, the processing of information pertaining to those rectangles is facilitated. Something more is true, which is that the rectangles themselves define the location to which attention is allocated.

It is in this sense that Norman, Heywood and Kentridge's (2013) experiment shows that attention can be paid to rectangles that are completely invisible. The participants in this experiment were sat in front of a screen that they consciously experienced as if it were just an undifferentiated grid of blurry, shimmering patches. Sixteen evenly distributed locations within this grid were empty. One of these empty locations would fill with a white disk for 160 ms. Another disk would

then appear, immediately after this attention-directing cue had disappeared. It would either be in the same location as the cue, or else in another of the empty cells in the grid. This second disk was either red or green. The participants pressed a button to indicate its colour. Their reaction times were measured.

The important trials in this paradigm are those in which the coloured disk appears in one of the uncued locations. By changing the orientations and locations of their unseen rectangles, Norman, Heywood and Kentridge (2013) could determine whether or not these uncued target locations did or did not fall within the same rectangle as the attention-directing cue, and they could do this without changing the locations of (or distances between) the targets and the cues. They could therefore arrange for the cue and the target to fall in the same or in different rectangles, without there being any consciously visible difference between those trials in which the cue and the target fell within the same rectangle and those in which they fell within different rectangles. And yet the effect of the attention-directing cue on the reaction times of their participants were different in these different cases. The attention-directing cues in Norman, Heywood and Kentridge's experiment therefore seem to draw attention to a region the boundaries of which are defined by the unseen rectangles. They do so in such a way as to facilitate the processing of items that occur within those boundaries. In such a case there is no sense in denying that attention is paid to the unseen rectangles. This is, therefore, a counter-example to the claim that one must be conscious of a thing in order to pay attention to it. It establishes the falsity of α.

3. Consequences of Rejecting α

Norman, Heywood and Kentridge's (2013) experiment is bad news for the claim that my 2008 paper took to be a natural one. It is also bad news for the theory of attention proposed by Smithies.

Smithies claims that 'Attention is what makes information fully accessible for use in the rational control of thought and action' (Smithies, 2011, p. 248). His article starts by taking this functional role to be essential to attention. The goal of the article is to derive a claim about attention's relation to consciousness from there. Smithies makes the claim just quoted as a claim about the functional role that attention has, essentially and fundamentally. He is not making a claim about one function that attention just happens to fulfil. The 'is' of his claim is intended as the 'is' of identity. But the attention elicited by the cue in Norman, Heywood and Kentridge's experiment shows this identity

claim to be false. Contrary to what Smithies' claim would entail, the participants in Norman, Heywood and Kentridge's experiment have no rational access to the rectangles to which they are attending. If it happened that there was something about the experience of these experimental stimuli that prompted the subjects to form a belief about rectangles having been present on the screen, that belief could only be a hunch or a guess. The attention that is paid to these rectangles cannot make any beliefs that are formed about them into rational ones. Nor does attention to the rectangles make the information about those rectangles 'fully-accessible'. The rectangles remain invisible, even when attention is being paid to them. Smithies' rational-access view of attention is therefore refuted by the Norman, Heywood and Kentridge (2013) experiment.

The news is almost as bad for Watzl. Since attention can be paid to the rectangles in Norman, Heywood and Kentridge's experiment, and since these do not feature in the contents of consciousness (and so cannot possibly participate in any conscious peripherality relation), Watzl's (2011) theory of the way in which attention imposes a peripherality structure on the contents of consciousness cannot tell us about the essence of attention. It gets no grip on the attention paid by Norman, Heywood and Kentridge's participants, and so cannot be a complete account of what attention *is*.

The problem here is not just that Norman, Heywood and Kentridge's experiment reveals an awkward counter-example, which these philosophical theories might be adapted to accommodate. The problem, as we have already indicated, is that the result of their experiment reveals the explanatory project that these theorists are embarked on to be a misbegotten one. The possibility of attention being paid to unseen objects is a problem for *any* theory that attempts to account for attention by specifying the circumstances in which consciousness comes to have an attended foreground. Any variation on the themes that Watzl and Smithies have established will be vulnerable to some version of this same problem. If there is attention without consciousness, as the Norman, Heywood and Kentridge (2013) experiment indicates, then attending cannot be any particular way of being conscious, however carefully the way in question might be specified.

We started by outlining two explanatory strategies, no more than one of which could be correct: Smithies and Watzl pursued the consciousness-first strategy in their explanations of attention; Prinz pursued the attention-first strategy in his explanation of consciousness. We have now seen that, so far as the debate about our claim α is concerned, it is

Watzl and Smithies whose explanatory project seems to be in trouble. Prinz's explanatory strategy is not threatened by the falsity of α. Prinz's own way of pursuing that strategy does, however, face a problem.

This problem arises because Prinz's attention-based explanation of consciousness is one in which attention is understood to essentially involve the projection of sensory information into working memory. He is quite explicit about this, writing that: 'attention can be identified with the processes that allow information to be encoded in working memory' (Prinz, 2012, p. 93). The Norman, Heywood and Kentridge (2013) experiment undermines this identification.

The participants in that experiment are, as we have said, paying attention to one or other of the invisible rectangles. If Prinz's account of attention were correct then a process would need to be taking place that allows information about that particular rectangle to be encoded in working memory. But there is no reason to suppose that information about the rectangles is stored in any sort of memory at all. There is no reason to suppose that it is stored beyond the time when the rectangles are perceptually present. Nor is there any evidence that any sort of *executive* 'work' is devoted to the information pertaining to one of these rectangles in particular: the subjects in Norman, Heywood and Kentridge's experiment are not even disposed to guess that the rectangles are present. When they do make guesses about the presence of these rectangles, there is *no* information about the rectangles in the guesses that they make (*ibid.*, Fig. 2b). There is therefore no indication that any process is taking place that allows information about one of these rectangles to be encoded in working memory. The attention that is paid in Norman, Heywood and Kentridge's study seems to be an instance of attention that has no connection to working memory at all.

Norman, Heywood and Kentridge's findings do not require the abandonment of Prinz's theory when that theory is being understood — without reference to attention — as a theory of the way in which consciousness is related to the information processing taking place in the brain. Prinz can retain the central claim of his AIR theory, which is that consciousness occurs when 'intermediate-level' representations are sent to working memory (see Prinz, 2012, p. 293). The Norman, Heywood and Kentridge experiment gives no counter-example to that, since the unseen objects in that experiment are defined by boundaries of visual texture, and so are not objects that have been represented at the *intermediate* level, in Prinz's sense of 'intermediate'. Prinz's theory of *consciousness* is not threatened by the Norman,

Heywood and Kentridge (2013) finding, but his theory of attention is. That finding gives a counter-example to Prinz's claim that 'attention can be identified with the processes that allow information to be encoded in working memory', and so undermines Prinz's suggestion that, in giving his 'AIR' theory, he is giving an *attention*-based theory of consciousness.

This is only a superficial blow to Prinz's project. Prinz advertises the AIR theory as an account of 'how *attention* engenders experience', but he has suggested that he would be willing to give up the word 'attention', if it misrepresents his working-memory based theory of consciousness (Prinz, 2012, p. 95). The Norman, Heywood and Kentridge (2013) finding shows that he would be right to do so.

4. Attention and Agency

If the prospects are bad for Prinz-style theories, in which consciousness is explained with reference to attention, and for Smithies-and-Watzl-style theories, in which attention is explained with reference to consciousness, then the way is open for a theory in which these two phenomena are taken to be independent. Such a theory can still be made compatible with the idea that attention is intimately related to processes in which consciousness is typically involved. It need only maintain that this relation to consciousness is owing to a contingent feature of our psychology. My own work gives such a theory of attention (Mole, 2011). The attention/consciousness link is also taken to be inessential in the selection for action theory of attention, which has recently been reanimated in the work of Wayne Wu (2014; 2011; Allport, 1987; Neumann, 1987).

In these theories, attention is seen as being essentially connected to the exercise of agency, in a broad sense of 'agency', which is applicable to any creature that has and uses a will. According to Wu's theory, a subject is paying attention just in case she is *navigating* through the space of possible bodily and mental actions that are currently available to her, rather than drifting through that space at random. According to my own theory — in which attention is identified with 'cognitive unison' — an agent is doing something attentively only if she is doing something under the guidance of her understanding of that thing (and only if the resources that she can bring bear, as part of her current attempt to do the thing, are not occupied with processing of a sort that is irrelevant to her attempt). In neither theory is attention primarily a phenomenon of conscious experience. Attention involves experience only to the extent that the world we experience is a world

in which we act (where, again, 'action' is construed broadly, and includes such epistemic actions as inspecting, surveying, or reasoning).

There are some immediate advantages to this breaking of the connection between attention and consciousness.

The first is that it allows for the literal truth of sentences in which attention is attributed to unconscious entities, as when we say that the Manchester United football team paid more attention to defence in the first half of the game, and more attention to attack in the second, or that it paid more attention to the FA Cup in the 1980s, and more attention to the Champion's League in the 1990s. It is plausible that such claims require us to think of that whole football team as being an *agent*, capable of doing things (such as winning leagues) that are not just the several actions of its several members. Taking such claims to be literally true does *not* require us to think — what would be much more problematic — that the team constitutes a single emergent subject of conscious experiences, additional to the experiences of its constituent members. That is an indication that our attributions of attention attribute a phenomenon of agency, and not of consciousness.

A further advantage of breaking the link between attention and consciousness is that it allows for there to be true claims about attention over timescales where the attribution of any particular phenomenology would seem to be misplaced. We may want to allow that it can be literally true that a person started attending to her career in her twenties, having been attending to other things in her teens. Such a claim would tell us about the way in which this agent structured the execution of her projects over these decades, not about any decade-long tendencies in her phenomenal experience. Again the facts about attention seem most fundamentally to be facts about *agency*, not facts about conscious experience (although they do, of course, have implications that relate to conscious experiences, when the agent in question is a conscious one).

An anonymous referee for this journal complains that 'the use of the term "attention" with respect to the behaviour of football teams or career advancement is very different from its use elsewhere'. This referee therefore takes the arguments of the previous two paragraphs as spurious. It is true that these examples have a different character from the previous ones, but I do not think that they are cases in which the word 'attention' is used with a different sense. To introduce these examples to the discussion is not to change the subject (as it would be if I were suggesting that our theory of attention should apply not only to the focusing of mental resources but also to the posture of the

soldier who is about to present arms). I suppose — as an axiom of philosophical methodology — that our theory of attention should, if possible, be a theory of the phenomenon referred to by the English word 'attention'. If that is right then our theory should apply to the full range of cases. To broaden our diet of examples is not to change the subject. The anonymous referee is no doubt right that the psychological processes that are involved in these last examples are different from the psychological processes that are involved in the 200 ms-long instances of attention that are studied in typical laboratory experiments, including those that we have been considering above, but it would be a mistake to suppose that, since different *processes* are involved in these different instances of 'attention', the word 'attention' must be being used with a different sense when we refer to them. To suppose that would be to beg the larger question (considered at length in Mole, 2011), by supposing that 'attention' is a process-first term, and not an adverbial one.

Attention has not, traditionally, enjoyed a place alongside consciousness, free will, or intentionality, on the list of those features of the mind that resist scientific explanation, and that therefore demand philosophical scrutiny. It is partly as a result of this that the theory of attention has become one of the most thoroughly achieved theories in cognitive psychology's picture of the mind. It is understandable that philosophers have hoped that theory could be put to philosophical work. And — it being the received wisdom that 'Consciousness is what makes the mind body problem really intractable' (Nagel, 1974) — it is understandable that these philosophers have hoped the theory of attention would be related in some way to the theory of consciousness. The above considerations suggest that that hope is mistaken. The theory of attention may indeed have philosophical work to do, but attention is more intimately related to the mysteries of agency than to the mysteries of consciousness. Getting clear on the proper relationships between these mysteries — on how many of them there are, and on which are more basic than the others — will be an essential step towards their eventual solution.[1]

References

Allport, D.A. (1987) Selection-for-action: Some behavioural and neurophysiological considerations of attention and action, in Heuer, H. & Sanders,

[1] Thanks to Kathleen Akins and Robert Kentridge, for useful discussions, and to Liam Norman, for providing Figure 1.

A.F. (eds.) *Perspectives on Perception and Action*, pp. 395–419, Hillsdale, NJ: Erlbaum.
Barbur, J.L., Ruddock, K. & Waterfield, V. (1980) Human visual responses in the absence of the geniculo-calcarine projection, *Brain*, **103** (4), pp. 905–928.
de Brigard, F. (2010) Consciousness, attention, and commonsense, *Journal of Consciousness Studies*, **17** (9–10), pp. 189–201.
Egly, R., Driver, J. & Rafal, R.D. (1994) Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects, *Journal of Experimental Psychology: General*, **123** (2), pp. 161–177.
Kentridge, R. & Heywood, C. (2001) Attention and alerting: Cognitive processes spared in blindsight, in De Gelder, B., De Haan, E. & Heywood, C. (eds.) *Out of Mind: Varieties of Unconscious Processes*, pp. 163–181, New York: Oxford University Press.
Kentridge, R.W., Heywood, C.A. & Weiskrantz, L. (2004) Spatial attention speeds discrimination without awareness in blindsight, *Neuropsychologia*, **42** (6), pp. 831–835.
Kentridge, R.W., de-Wit, L.H. & Heywood, C.A. (2008) What is attended in spatial attention?, *Journal of Consciousness Studies*, **15** (4), pp. 105–111.
Kentridge, R.W., Nijboer, T.C. & Heywood, C.A. (2008) Attended but unseen: Visual attention is not sufficient for visual awareness, *Neuropsychologia*, **46**, pp. 864–869.
Koch, C. & Tsuchiya, N. (2008) Response to Mole: Subjects can attend to completely invisible objects, *Trends in Cognitive Sciences*, **12** (2), pp. 44–45.
Mole, C. (2005) *Attention is Cognitive Unison*, PhD Dissertation, Princeton University.
Mole, C. (2008a) Attention and consciousness, *Journal of Consciousness Studies*, **15** (4), pp. 86–104.
Mole, C. (2008b) Attention in the absence of consciousness?, *Trends in Cognitive Sciences*, **14** (2), p. 44.
Mole, C. (2011) *Attention is Cognitive Unison*, New York: Oxford University Press.
Nagel, T. (1974) What is it like to be a bat?, *The Philosophical Review*, **83** (4), pp. 435–450.
Neumann, O. (1987) Beyond capacity: A functional view of attention, in Heuer, H. & Sanders, A.F. (eds.) *Perspectives on Perception and Action*, pp. 361–394, Hillsdale, NJ: Erlbaum.
Norman, L.J., Heywood, C.A. & Kentridge, R.W. (2013) Object-based attention without awareness, *Psychological Science*, **24** (5), pp. 836–843.
Prinz, J.J. (2010) When is perception conscious?, in Nanay, B. (ed.) *Perceiving the World: New Essays on Perception*, pp. 310–332, Oxford: Oxford University Press.
Prinz, J.J. (2011) Is attention necessary and sufficient for consciousness, in Mole, C., Smithies, D. & Wu, W. (eds.) *Attention: Philosophical and Psychological Essays*, pp. 174–203, New York: Oxford University Press.
Prinz, J.J. (2012) *The Conscious Brain*, New York: Oxford University Press.
Rosenthal, D. (1986) Two concepts of consciousness, *Philosophical Studies*, **49** (3), pp. 329–359.
Smithies, D. (2011) Attention is rational access consciousness, in Mole, C., Smithies, D. & Wu, W. (eds.) *Attention: Philosophical and Psychological Essays*, pp. 247–273, New York: Oxford University Press.
van Marle, K. & Scholl, B.J. (2003) Attentive tracking of objects versus substances, *Psychological Science*, **14** (5), pp. 498–504.

Watzl, S. (2011) Attention as structuring of the stream of consciousness, in Mole, C., Smithies, D. & Wu, W. (eds.) *Attention: Philosophical and Psychological Essays*, pp. 145–173, New York: Oxford University Press.
White, A.R. (1964) *Attention*, Oxford: Basil Blackwell.
Wu, W. (2011) Attention as selection for action, in Mole, C., Smithies, D. & Wu, W. (eds.) *Attention: Philosophical and Psychological Essays*, pp. 97–116, New York: Oxford University Press.
Wu, W. (2014) *Attention*, Oxford: Routledge.

Paper received February 2014; revised June 2014.