

Goldsmiths Research Online

*Goldsmiths Research Online (GRO)
is the institutional research repository for
Goldsmiths, University of London*

Citation

Velmans, Max. 2007. A natural account of phenomenal consciousness. *NeuroQuantology*, 3(3), pp. 164-179. [Article]

Persistent URL

<https://research.gold.ac.uk/id/eprint/26184/>

Versions

The version presented here may differ from the published, performed or presented work. Please go to the persistent GRO record above for more information.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Goldsmiths, University of London via the following email address: gro@gold.ac.uk.

The item will be removed from the repository while any claim is being investigated. For more information, please contact the GRO team: gro@gold.ac.uk

A Natural Account of Phenomenal Consciousness

Max Velmans

Abstract

Physicalists commonly argue that conscious experiences are nothing more than states of the brain, and that conscious qualia are observer-independent, physical properties of the external world. Although this assumes the 'mantle of science,' it routinely ignores the findings of science, for example in sensory physiology, perception, psychophysics, neuropsychology and comparative psychology. Consequently, although physicalism aims to naturalise consciousness, it gives an unnatural account of it. It is possible, however, to develop a natural, nonreductive, reflexive model of how consciousness relates to the brain and the physical world. This paper introduces such a model and how it construes the nature of conscious experience. Within this model the physical world as perceived (the phenomenal world) is viewed as part of conscious experience not apart from it. While in everyday life we treat this phenomenal world as if it is the "physical world", it is really just one biologically useful representation of what the world is like that may differ in many respects from the world described by physics. How the world as perceived relates to the world as described by physics can be investigated by normal science (e.g. through the study of sensory physiology, psychophysics and so on). This model of consciousness appears to be consistent with both third-person evidence of how the brain works and with first-person evidence of what it is like to have a given experience. According to the reflexive model, conscious experiences are really how they seem.

Key Words: Phenomenal consciousness, consciousness, reflexive model, qualia, neuropsychology

NeuroQuantology 2005; 3 :164-179

DEFINING CONSCIOUSNESS

There are many differences of opinion about how to define consciousness. This uncertainty about how to define consciousness is partly brought about by the way global theories about consciousness (or even about the nature of the universe) have intruded into definitions. For example, "substance dualists" such as Plato, Descartes, and Eccles believe the universe to consist of two fundamental kinds of stuff, material stuff and the stuff of consciousness (a substance associated with soul or spirit). "Property dualists" such as Sperry and Libet take consciousness to be a special kind of property that is itself nonphysical, but which emerges from physical systems such as the brain once they attain a

certain level of complexity. By contrast, "reductionists" such as Crick (1994) and Dennett (1991), believe consciousness to be nothing more than a state or function of the brain. Within cognitive psychology, there are many proposals which identify consciousness with some aspect of human information processing, for example with working memory, focal attention, a central executive, and so on.

Fortunately, definitions need not be final for research to get under way. It is enough that, for given investigative purposes, definitions are sufficiently similar for different investigators to be able to agree that they are investigating the same thing. As science begins to unravel the causes of consciousness, the functions of consciousness, how consciousness relates to nonconscious processing in the brain and so on, our understanding of what consciousness is will deepen - for the reason that such relationships form part of the meaning of the term (its connotative meaning, or sense). Such mutual focusing of attention followed by exploration of the nature of what is attended to (and how it relates to other things) is fundamental to how phenomena come to be understood in a socially shared way. In this respect, coming to understand the nature of consciousness is no different to coming to understand the nature of anything else.

Nevertheless, before any investigation can begin, one has to "point to" or "pick out" the phenomena to which the term refers and, by implication, what is excluded. In everyday life there are two contrasting situations which inform our understanding of the term "consciousness". We have knowledge of what it is like to be conscious (when we are awake) as opposed to not being conscious (when in dreamless sleep). We also understand what it is like to be conscious of something (when awake or dreaming) as opposed to not being conscious of that thing. This everyday understanding provides a simple place to start. A person, or other entity, is conscious if they experience something; conversely, if a person or entity experiences nothing they are not conscious. Elaborating slightly, we can say that when consciousness is present, phenomenal content is present. Conversely, when phenomenal content is absent, consciousness is absent. This stays very close to everyday usage and, for this paper, it is all that we need. To minimise confusion, I will also stay as close as possible to everyday, natural language usage for related terms. In common usage, the term "consciousness" is often synonymous with "awareness" or "conscious awareness." Consequently, I will use these terms interchangeably." The "contents of consciousness" encompass all that we are conscious of, aware of, or experience. These include not only experiences that we commonly associate with ourselves, such as thoughts, feelings, images, dreams, body sensations and so on, but also the experienced three-dimensional world (the phenomenal world) beyond the body surface.

Of course, to learn what something is, it is useful in the initial instance to know where it is, so that one can point to it - enabling the attention of different investigators to be focused upon it. But where does one point, when one is pointing at phenomenal consciousness?

Where dualists and reductionists think consciousness to be
According to Descartes the material world is composed of *res extensa*, a substance that has both location and extension in space. Consciousness is formed out of *res cogitans* a substance which thinks, but which has no location or extension in space. If this is right, then

one cannot point at consciousness, as it has no location. At best, one might be able to point at the place where consciousness interfaces with the material world. According to Descartes this is at the pineal gland located in the centre of the brain. Physicalist and functionalist philosophers (e.g. Searle 1992; Dennett 1991) argue that consciousness is nothing more than a state or function of the brain. It might be difficult to point with any precision at such states or functions as they are likely to be distributed properties of large neuronal populations (cf Dennett & Kinsbourne 1992). Nevertheless, if one had to point one would point at the brain. In short, classical dualists and reductionists disagree vehemently about what conscious is, but they agree (roughly) about where it is. In so far as consciousness can be located at all, that location is somewhere in the brain.

A common-sense view of conscious phenomenology

In Velmans (1990, 2000) I have argued that this currently popular view has no basis either in science or in everyday experience. In order to decide where consciousness is (or whether it has any location) one has to attend to its actual phenomenology. It is true that there are some experiences which seem to be poorly localised in space, or at best localised somewhere in the head or brain, just as dualists and reductionists claim. Examples include thoughts and vague feelings such as the verbal thoughts, feelings of understanding and so on that accompany the reading of this text. However, most experiences have a very different phenomenology, for example experiences of the body or of the external world.

Let me illustrate with a very simple example. Suppose you stick a pin in your finger and experience a sharp pain. Within philosophy of mind pain is generally regarded as a paradigm case of a conscious, mental event. But where is the pain? Hampered by their theoretical presuppositions, dualists and reductionists take this to be a rather difficult question. However, if forced to point they would point (vaguely) in the direction of the brain (see comments by Nagel, Harnad, Searle, Marcel, and Dennett, following Velmans 1993). I take this to be a very simple question. The pain one experiences is in the finger. If one had to point at it one should point at where the pin went in. Any reader in doubt on this issue might like to try it.

Let me be clear that this sharp difference of opinion is about the experienced pain and not about the antecedent physical causes (the deformation and damage to the skin produced by the pin) or about the neural causes and correlates of pain. The proximal neural causes and correlates of pain are undoubtedly located in the brain. But the neural causes and correlates of a given experience are not themselves that experience. In science, causes and correlates are not ontological identities. I have given a detailed analysis of how causes and correlates relate to ontological identities in Velmans (1998, 2000), so I won't labour the point here.

This subjective location of pains in parts of the body rather than "nowhere" or "in the brain" exemplifies a general principle that leads one away from both dualism and reductionism towards a "reflexive" model of how consciousness relates to the brain and the physical world (cf Velmans 1990). In many respects, there is no difference between these theoretical positions. For example, dualism, reductionism and the reflexive model agree that there are physical and neurophysiological causes and correlates of a given experience within the brain - and that we can leave it to science to discover what these are. But they disagree

about the nature and location of the effects (the resulting experiences). Dualists claim that, being constructed out of *res cogitans*, experiences have no location or extension in space (although they interface with the brain). Reductionists claim that, being brain states or functions, all experiences must be in the brain (in spite of how they seem). According to the reflexive model, the only evidence about conscious phenomenology comes from first-person sources. Consequently, the properties of that phenomenology can only be determined from first-person sources. For conscious appearances, the appearance is the reality (Searle, 1992). Consequently, if a pain appears to be in the finger, then that is where the pain is. The damage produced by a pin in the finger, once it is processed by the brain, winds up as a phenomenal pain in the finger, located more or less where the pin went in. That is why the entire process is called "reflexive".

Notice that if one stabs one's finger with a pin, and one attends to the consequent pain phenomenology, one has no additional, experience of pain either "nowhere" or in the brain. Nor can any phenomenal pain "nowhere" or in the brain be observed by an external observer (from a third-person perspective only its neural causes and correlates can be observed). Given that there is no first- or third-person evidence for phenomenal pain "nowhere" or in the brain, I suggest that this is a theoretical fiction, introduced by dualist and reductionist thinkers in order to make their models work. Only the reflexive model is consistent with the evidence of common sense.

To put the basic principle in a more general way: experiences are where we experience them to be. Figure 1, for example, illustrates a similar process with a phenomenal cat. As before, some entity or event innervates sense organs and initiates perceptual processing, although in this case the initiating entity is located beyond the body surface in the external world. As before, afferent neurons, and cortical projection areas are activated, along with association areas, long-term memory traces and so on, and neural representations of the initiating event are eventually formed within the brain - in this case, neural representations of a cat. But the entire causal sequence does not end there. The subject S also has a visual experience of a cat and, as before, we can ask what this experience is like. In this case, the proper question to ask is, "What do you see?" According to dualism, S has a visual experience of a cat "in her mind". According to reductionists there seems to be a phenomenal cat "in S's mind" but this is really nothing more than a state of her brain. According to the reflexive model, while S is gazing at the cat, her only visual experience of the cat, is the cat she sees out in the world. If she is asked to point to this phenomenal cat (her "cat experience"), she should point not to her brain but to the cat as-perceived, out in space beyond the body surface. In this, S is no different from an external observer E. The cat as perceived by S is the same cat as perceived by E (albeit viewed from S's perspective rather than from E's perspective). That is, an entity in the world is reflexively experienced to be an entity in the world.

Of course, not all the entities and events we experience have such a clear location and extension in three-dimensional phenomenal space. We also have "inner" experiences such as verbal thoughts, images, feelings of knowing, experienced desires and so on. Such inner experiences really do seem to have a phenomenology of the kind that characterise Descartes' *res cogitans*. One might argue that verbal thoughts have a rough location, in that they seem to be "in the head" (in the form of inner speech) rather than in one's foot, or

free-floating out in space, but they are not clearly located in the manner of pains and cats. However, the reflexive process is the same. The cognitive processes which give rise to thoughts, feelings of knowing and so on originate in the mind/brain, although these processes are unlikely to have a precise location in so far as they engage the mass action of large, distributed, neuronal populations. Consequently, in so far as these processes are experienced, they are reflexively experienced to be roughly where they are (in the head or brain).

There is far more to be said about conscious phenomenology and its relation to the brain and physical world. But, if I am right so far, even a cursory examination of what we actually experience poses a fundamental challenge to dualist and reductionist presuppositions about what it is that they need to explain. Both dualism and reductionism assume experiences to be quite different from the perceived body and the perceived external world (perceived bodies and worlds are out-there in space, while experiences of bodies and worlds are "nowhere" or in the brain). But the reflexive model suggests that in terms of phenomenology there is no actual separation between the perceived body and experiences of the body or between the perceived external world and experiences of that world. It goes without saying that when one has a conscious thought, there isn't some additional experience of a thought "in the mind". But neither is there a phenomenal pain "in the mind" (without location and extension) in addition to the pain one experiences in the finger if one stabs it with a pin. And there isn't a phenomenal cat "in the mind" in addition to the cat one sees out in the world. Applying Occam's razor, the reflexive model gets rid of them.




Figure 1. A Reflexive Model of Perception

But the reflexive model does not get rid of conscious phenomenology. Thoughts, pains and phenomenal cats are experienced to have very different "qualia" (along with different locations and extensions), but they are nevertheless aspects of what we experience. Together, such inner experiences, bodily sensations, and external experienced entities and events comprise the contents of our consciousness - which are none other than our everyday phenomenal world.

Who else says this?

To those immersed in dualist or reductionist modes of thought this proposed expansion of the contents of consciousness to include the entire phenomenal world may seem radical and the notion that many experiences have a precise location and extension might appear strange. But, thus far, this proposal is hardly new. In one or another form it appears in the work of George Berkeley, Immanuel Kant, C.H. Lewes, W.K. Clifford, Ernst Mach, Morton Prince, William James, A.N. Whitehead, Charles Sherrington, Bertrand Russell, R. Brain, Wolfgang Köhler and Karl Pribram. Similar analyses of what consciousness seems to be like have also recently been given by Antti Revonsuo and Michael Tye.

William James (1904) for example, suggests that to convince oneself about where experiences are the observer only needs to

"... begin with a perceptual experience, the 'presentation', so called, of a physical object, his actual field of vision, the room he sits in, with the book he is reading as its centre, and let him for the present treat this complex object in the commonsense way as being 'really' what it seems to be, namely, a collection of physical things cut out from an environing world of other physical things with which these physical things have actual or potential relations. Now at the same time it is just those self-same things which his mind, as we say, perceives, and the whole philosophy of perception from Democritus's time downwards has been just one long wrangle over the paradox that what is evidently one reality should be in two places at once, both in outer space and in a person's mind. 'Representative' theories of perception avoid the logical paradox, but on the other hand they violate the reader's sense of life which knows no intervening mental image but seems to see the room and the book immediately just as they physically exist".

One insight, of course, does not make a theory. While the philosophers and scientists mentioned above agree that some experiences appear to have location and spatial extension, there is widespread disagreement about what this implies about the nature of consciousness and its relation to the physical world. Berkeley for example is an idealist, James a neutral monist, Whitehead a process theorist, and Tye, a physicalist. In Velmans (2000), I develop reflexive monism, an analysis of what is going on which is none of these (although it incorporates elements of many positions).

It is only possible to introduce a few consequences of the reflexive model and how this translates into reflexive monism (a broad philosophical position) in a brief journal article. So, by way of introduction, I will focus here on just one fundamental issue: how it makes sense of conscious intentionality (that consciousness is consciousness of something). To set the theory in a philosophical context I will also contrast the analysis I develop with that of Armstrong (1968), Block (1997) and Tye (1995).

What do experiences represent?

In dualism and reductionism it is easy to see what experiences of the external world represent. Percepts of objects "in the mind" or "in the brain" represent the objects we see out in the world. But, at first glance, the intentionality of conscious experiences would seem to be a problem for the reflexive model. If experiences of objects and objects as-perceived are phenomenologically identical, as argued above, then what do experiences of objects represent? One might ask the same question about the experienced body and about "inner" experiences.

And there is a related question. According to the reflexive model, what we commonly refer to as the "physical world" is just the world we experience. However, this clearly remains very different to the world described by modern physics (the world of quantum mechanics, relativity theory, grand unified theory and so on). So how does the phenomenal, "physical world" relate to the world described by physics?

A reflexive model of how consciousness relates to the brain and the physical world

The reflexive model shown in Figure 1 suggests that all experiences result from a reflexive interaction of an observer with an observed. For the purposes of illustrating how this interaction works to produce different kinds of experience, these can be subdivided into three categories:

- (1) experiences of the external world (which seem to have location and extension)
- (2) experiences of the body (which seem to have location and extension)
- (3) "inner" experiences (thoughts, images, feelings of knowing and so on) which have no clear location and extension in phenomenal space, although they can be loosely said to be "in the head or brain".

Figure 1 illustrates one example of a reflexive interaction resulting in an experience (a visual percept) of a phenomenal cat. In this case, the initiating stimulus (the observed) is an entity located in space beyond the body surface that interacts with the visual system of the observer to produce an experienced entity out in space beyond the body surface. As noted above, a similar reflexive interaction takes place when the initiating stimulus is on the surface of (or within) the body, or within the brain itself to produce experienced entities and events on the surface of (or within) the body, or in the head or brain itself.

What is going on? Following current conventions in the psychology of perception, I assume that the brain constructs a "representation" or "mental model" of what is happening, based on the input from the initiating stimulus, expectations, traces of prior, related stimuli stored in long-term memory, and so on (cf Rock, 1997). Such mental models encode information about the entities and events that they represent in formats determined by the sensory modality that they employ. Visual representations of a cat, for example, include encodings for shape, location and extension, movement, surface texture, colour, and so on. In addition, I suggest that the way information (in a given mental model) appears to be formatted depends on the observational arrangements. The information appears in different forms to the subject (S) and an external observer (E) for the reason that the means available to S and E for accessing the information in that mental model differ (cf Velmans, 1991).

An external observer, inspecting a subject's brain, has to rely on his own exteroceptive systems (typically vision) aided by physical equipment (PET scans, fMRI and so on). Viewed in

this way (from this third-person perspective) a visual mental model in the subject's brain might appear in the form of neural activation in a series of relatively distinct feature maps distributed throughout the subject's visual system. We do not know precisely what is required to make such neural representations conscious. However, given the integrated nature of visual experiences, it is reasonable to assume that when such distributed neural activities do become conscious they must be bound together in some way, perhaps through synchronous 40 Hz oscillations. We may also expect there to be observable (physical) influences on the pattern of activity embodied in the mental model from existing memory traces (corresponding to the effects of expectation, stored knowledge and so on). Whatever the fine detail turns out to be like, viewed from E's perspective, the information (about the cat) in S's mental model is likely to take a neural, or other physical form. In terms of what E can directly observe of S's mental model, that is the end of the scientific story.

However, the observational arrangement by which the subject accesses the information in her own mental model is entirely different. As with E, the information in her own mental model is translated into something that she can observe or experience - but all she experiences is a phenomenal cat out in the world. While she focuses her attention on the cat she does not become conscious of having a "mental model of a cat" in the form of neural states. Nor does she have an experience of a cat "in her head or brain". Rather, she become conscious of what the neural states represent - an entity out in the external world. In short, the information encoded in S's mental model (about the entity in the world) is identical whether viewed by S or by E, but the way the information appears to be formatted depends on the perspective from which it is viewed.

Let me illustrate with a simple analogy. Let us suppose that the information encoded in the subject's brain is formed into a kind of neural "projection hologram." A projection hologram has the interesting property that the three-dimensional image it encodes is perceived to be out in space, in front of its two-dimensional surface, provided that it is viewed from an appropriate (frontal) perspective and it is illuminated by an appropriate (frontal) source of light. Viewed from any other perspective (from the side or from behind) the only information one can detect about the object is in the complex interference patterns encoded on the holographic plate. In analogous fashion, the information in the neural "projection hologram" is displayed as a visual, three-dimensional object out in space only when it is viewed from the appropriate, first-person perspective of the perceiving subject. And this happens only when the necessary and sufficient conditions for consciousness are satisfied (when there is "illumination by an appropriate source of light"). Viewed from any other, external perspective the information in S's "hologram" appears to be nothing more than neural representations in the brain (interference patterns on the plate).

The "projection hologram" is, of course, only an analogy- but it is useful in that it shares some of the apparently puzzling features of conscious experiences. The information displayed in the three-dimensional holographic image is encoded in two-dimensional patterns on a plate, but there is no sense in which the three-dimensional image is itself "in the plate". Likewise, there is no sense in which the phenomenal cat observed by S is "in her head or brain." In fact, the 3-D holographic image does not even exist (as an image) without an appropriately placed observer and an appropriate source of light. Likewise, the existence of the phenomenal cat requires the participation of S, the experiencing agent, and all the

conditions required for conscious experience (in her mind/brain) have to be satisfied. Finally, a given holographic image only exists for a given observer, and can only be said to be located and extended where that observer perceives it to be! S's phenomenal cat is similarly private and subjective. If she perceives it to be out in phenomenal space beyond the body surface, then, from her perspective, it is out in phenomenal space beyond the body surface.

Perceptual projection

Unconscious mind/brain processes construct experienced realities in which our phenomenal heads appear to be enclosed within three-dimensional, phenomenal worlds, not the other way around. But the mental models that encode information about these 3D experienced realities are "in the head or brain". Given this, how do phenomenal cats and other phenomenal objects that are perceived to be located and extended in space get to be out there? It is clear that nothing physical is projected by the brain. There are for example no light rays projected through the eyes to illuminate the world, contrary to the beliefs of ancient Greek thinkers such as Empedocles (cf Zajonc, 1993). Rather, "perceptual projection" is a psychological effect produced by unconscious perceptual processing. The projection hologram has a number of features that might be usefully incorporated into a causal explanation of such effects, but it is not intended to be a literal theory of what is taking place in the mind/brain. Right now, we just don't know how it is done. Of course, not fully understanding *how* it happens, does not alter the fact *that* it happens - and the experimental and clinical evidence for perceptual projection is considerable. I have reviewed this elsewhere (in Velmans, 1990, 2000) and will not repeat that review here. Clinical and experimental examples include phantom limbs, hallucinations and virtual realities. A particularly striking example is reported by the neurologist Peter Brugger (1994) in a clinical case history of a 17 year-old man suffering from epilepsy caused by a lesion in his left temporal lobe. He was being treated with anti-convulsant drugs to control the condition and was scheduled for surgery when he experienced an "heautosopic" episode (a visual hallucination of his body combined with an out-of-body experience) which was disturbing in the extreme:

"The heautosopic episode, which is of special interest to the topic of this report, occurred shortly before admission. The patient stopped his phenytoin medication, drank several glasses of beer, stayed in bed the whole of the next day, and in the evening he was found mumbling and confused below an almost completely destroyed large bush just under the window of his room on the third floor. At the local hospital, thoracic and pelvic contusions were noted. The patient gave the following account of the episode: on the respective morning he got up with a dizzy feeling. Turning around, he found himself still lying in bed. He became angry about "this guy who I knew was myself and who would not get up and thus risked being late for work". He tried to wake the body in bed first by shouting at it; then by trying to shake it and then repeatedly jumping on his alter ego in the bed. The lying body showed no reaction. Only then did the patient begin to be puzzled about his double existence and become more and more scared by the fact that he could no longer tell which of the two he really was. Several times his body awareness switched from the one standing upright to the one still lying in bed; when in the lying bed mode he felt quite awake but completely paralysed and scared by the figure of himself bending over and beating him. His

only intention was to become one person again and, looking out of the window (from where he could still see his body lying in bed), he suddenly decided to jump out "in order to stop the intolerable feeling of being divided in two". At the same time, he hoped that "this really desperate action would frighten the one in bed and thus urge him to merge with me again". The next thing he remembers is waking up in pain in the hospital." (Brugger, 1994, pp 838-839).

In short, this patient mistakenly judged the hallucinated body on the bed to be his real one and tried to get rid of his real body (which he judged to be the hallucination) in order to become unified again - a powerful example of the constructed, projected nature of the body as-experienced. But we do not really need such striking examples to demonstrate that there is something interesting going on that needs explanation. The simple fact that this WORD appears to be out here on this page (rather than in your brain) illustrates that the phenomenon is both ubiquitous and real.

The world as-perceived is part-of the contents of consciousness. Some initial principles that follow from the analysis above should now be clear. Within the reflexive model the physical world as-perceived is part of the contents of consciousness. The contents of consciousness are not in some separate place or space "in the mind or brain". That is, in terms of phenomenology no clear separation exists between what we normally think of as the "physical world", the "phenomenal world", and the "world as-perceived". That said, the everyday physical world as-perceived does have to be distinguished from the more abstract world described by physics (and other sciences). According to the reflexive model, the physical world as-perceived is just one, biologically useful representation of the world that science might describe in many alternative ways. But, with our eyes open, what we normally call the "physical world" just is what we experience. There is no additional experience of the world "in the mind or brain".

How does the phenomenal, "physical world" relate to the world described by physics? In Velmans (2000, ch7) I give a detailed review of how the mind/brain system translates the energies described by physics into a world-as-experienced. I will not repeat that review here. Suffice it to say that the data from physics, sensory physiology, perception and psychophysics makes it clear that the perceived world "models" only a selection of the events and energies that physics describes. There are electromagnetic energies of many kinds that permeate space and even penetrate our bodies, to which our eyes (and other sense organs) are blind. There are signals produced by animals and insects to which our ears are deaf. Each sensory system has its own limits of resolution. Changes in light intensity of less than around 5% or in sound intensity of less than around 20% are not perceived as changes. A change in sound frequency from 1000 Hz to 1005 Hz produces a just noticeable rise in pitch but not a change from 4000 Hz to 4005 Hz. A change in electromagnetic wavelength from 480 to 481 nanometers will produce a noticeable change in hue, but not a change from 550 to 551 nanometers. Our sense of smell and taste monitor, but tell us little of the chemistry of the substances we inhale and ingest. Sensation and perception are limited in their spatial resolution to detect events of a size and distance that are relevant to normal human action and survival – beyond this we need microscopes and telescopes. Our sensory systems are also structured to detect events of a given duration. Light bulbs, for example,

actually flash 50 times per second (the frequency of the A.C. mains voltage). However, this "flicker frequency" is faster than the visual system can resolve which makes the light seem continuous. By contrast, the movement of a flower out of the earth is too slow to see, so one needs time-lapse photography to 'see' the movement.

The data from comparative psychology, and zoology also suggests that the "physical reality" perceived by humans is only one of many possible perceived realities. The precise mix of sensory, perceptual, cognitive and social capacities in each species is unique. Human sensory and perceptual systems perform broadly similar functions to those of other animals. But the sensitivity of sense organs, the range of energies to which they are tuned, and the way information detected by the sensors is subject to perceptual processing vary considerably from species to species. Consequently, the "physical reality" that we perceive is actually a peculiarly human world.

PHYSICALIST ACCOUNTS OF THE LOCATION OF QUALIA

As far as I can judge, the above account of how observer-dependent, perceived phenomena represent an independently existing "reality" which natural science might describe in other ways is consistent both with science and with common sense. However the observer-dependence of qualia such as colour, smell, taste and so on has been strongly resisted by some physicalist philosophers of mind. Their resistance is a consequence of their commitment to physicalism. If qualia such as "redness" are, in their essence, observer-dependent experiences, then it is not easy to reduce such qualia to "objective" states of the brain, no matter how brain states are construed. Armstrong (1968), for example, acknowledges that unless one can exclude properties such as "redness" from perception he would have to abandon his entire reductive programme, which claims perception to be nothing more than the capacity to make certain discriminations. But "redness" undeniably exists, so Armstrong is forced into the view that redness is an observer-independent, physical property of certain physical objects. (Having excluded such qualia from perception there is nowhere else for them to go!)

According to the reflexive model, colour appears only once light waves (in the visible waveband) have been translated by the visual system into colour experiences. That is, objects are only red if (a) they reflect light with the appropriate wavelengths (around 700 nm) and (b) the visual system translates that electromagnetic energy into a red colour experience. Of these two conditions, (b) is the more important. That is, the visual system can produce a colour experience without being innervated by light in the 700 nm region (for example in dreams, vivid imagery, and hallucinations). But, without visual systems of the appropriate kind, light waves of 700 nm have no colour at all (colour as such is not an electromagnetic property). By contrast, Armstrong claims that objects are "red" whether or not there is anyone to perceive them. As van der Heijden et.al (1997) note (in their commentary on a similar position adopted by Block, 1995), such a view simply does not take the natural sciences seriously.

"That there are colours in the external world is a naive idea, unsupported by physics, biology, or psychology. Ultimately, it presupposes that the representation (the perceived colour) is represented (as a perceived colour). A perceptual system performs its proper function when it distinguishes the relevant things in the outer world. For vision, the

information about these relevant things is contained in the structure and composition of the light reflected by the outer world that enters the eyes. For distinguishing the relevant things in the external world, a unique and consistent representation of the corresponding distinctions in the light is all that is required." (Van der Heijden, et al, 1997, p158).

However, according to Block (1997), van der Heijden et.al are "wildly, unbelievably wrong. They say that we should give up the idea that a rose or anything else is ever red. The only redness, they say, is mental redness. But why not hold instead that roses are red rejecting colors in the mind? Why not construe talk of red in the mind as a misleading way of expressing the fact that P-conscious states represent the world as being red? And a representations of red need not itself be red (like the occurrences of the word "red" here)." (P165).

Block is, of course, right to point out that neural representations of red roses need not themselves be coloured. But no one claims that they are. What is claimed is that once a normal, human visual system is activated in an appropriate way, a visual experience of a red colour will result, irrespective of whether that colour corresponds to a physical property out in the world. Penfield & Rasmussen (1950), for example, demonstrated that direct microelectrode stimulation of the visual system resulted in visual experiences, stimulation of the temporal lobe in auditory experiences, stimulation of the somatosensory system in tactile experiences, and so on. Given that such visual, auditory, and tactile qualia can exist in the absence of the physical properties that they normally represent, it is not easy to see how they can be reduced to such physical properties.

A case for "red" and other qualia being observer-independent properties of the world rather than properties of experience has also recently been put by Tye (1995). Tye argues (as I do in Velmans, 1990) that all "qualia" are representational. He also agrees that qualia such as "redness" do not seem to be "in the mind or brain" but seem to be firmly attached to objects in the world. But he has an entirely different explanation of why the qualia seem to be out there. According to Tye, this results from perceptual experiences being transparent:

"Why is it that perceptual experiences are transparent? When you turn your gaze inward and try to focus your attention on intrinsic features of these experiences, why do you always seem to end up attending to what the experiences are of? Suppose you have a visual experience of a shiny, blood-soaked dagger. Whether, like Macbeth, you are hallucinating or whether you are seeing a real dagger, you experience redness and shininess as outside you, as covering the surface of a dagger. Now try to become aware of your experience itself, inside you, apart from its objects. Try to focus your attention on some intrinsic feature of the experience that distinguishes it from other experiences, something other than what it is an experience of. The task seems impossible: one's awareness seems always to slip through the experience to the redness and shininess, as instantiated together externally. In turning one's mind inward to attend to the experience, one seems to end up scrutinizing external features or properties." (p135).

To summarise, both physicalism and the reflexive model posit the existence of an observer-independent physical world, but they take a different view about how this relates to the phenomenal world (the perceived physical world). I treat the perceived "physical world" as part-of what we experience, and suggest that this experience is one, biologically

useful representation of what is really there. These representations may or may not correspond well to the world described by physics. For example, the perceived distance of experienced events represents measured distance quite well for close events, but breaks down completely for large distances (e.g. the perceived distance of the sun gives little indication of its measured distance). Perceived location, and distance result from preconscious, mental modelling processes involving "perceptual projection."

According to Tye, however, inner representational states are "transparent." That is, we "see through" our representations of colour, smell and so on to colours and smells as they really are out in the world. Tye bases his case partly on how things appear to us, and partly on evidence that perceived qualia really do correspond quite well to properties measured by Physics.

As Tye notes, "Certainly we do not experience colors as perceiver-relative. When, for example, a ripe tomato looks red to me, I experience redness all over the facing surface of the tomato. Each perceptible part of the surface looks red to me. None of these parts, in looking red look to me to have a perceiver-relative property. I do not experience any part of the surface as producing a certain sort of response in me or anyone else. On the contrary, I surely experience redness as intrinsic to it, just as I experience the shape of the surface as intrinsic to it." (p145). Given that we experience such colours as not being perceiver-relative, he regards the view that they are perceiver-relative as "just not credible" (p145).

Given that physicalism routinely denies the reliability of appearances as a guide to what experiences are really like, Tye rests his case on shaky ground. There are many obvious counterexamples. The colours of surfaces may seem to be observer-independent, but the colours of after-images do not. If one stares at a red spot for a few minutes, for example, one will experience a green after-image that projects onto any surface that the eye fixates. The apparent size of the after-image also increases as the judged distance of the surface increases. So, if apparent, observer-dependence is to be the criterion of what is "mental", after-images are surely mental. The observer-dependence of colour attached to surfaces in the world also becomes evident once the visual system no longer functions in the normal way. In cases of red-green colour blindness, for example, red can no longer be distinguished from green – and in cases of achromatopsia the entire world appears in shades of grey! More fundamentally, the reason that surfaces just appear coloured (without any conscious contribution on our part) is due to the fact that visual processing operates preconsciously. That is, once visual scenes appear in conscious experience, the binding of colour with shape, movement and so on has already taken place! Finally, it is important to note that variations in how things are experienced cannot be used to decide whether or not things are experienced!

Tye's second main argument relies on evidence that in some circumstances the qualia/physical property correspondence may be relatively invariant. Colours remain fairly similar for example when viewed outdoors, indoors (illuminated by incandescent lamps), or through sunglasses. Tye asks, "Why should this be? Surely the most straightforward answer is that the human visual system has, as one of its functions, to detect the real, objective colors of surfaces. Somehow, the visual system manages to ascertain what colors objects

really have, even though the only information immediately available to it concerns wavelengths." (p146) After a review of some of the relevant evidence, Tye concludes that

"Colors are objective, physical features of objects and surfaces. Our visual systems have evolved to detect a range of these features, but those to which we are particularly sensitive are indirectly dependent on facts about us. In particular there are three types of receptor in the retina, each of which responds to a particular waveband of light, and the spectral reflectances of surfaces at those wavebands (that is, their disposition to reflect a certain percentage of incident light within each of the three bands) together determine the colors we see. So the colors themselves may be identified with ordered triples of spectral reflectances. An account of the same general sort may be given for smells, tastes, sounds, and so on." (Tye, 1995, p150)

Tye is right to point out that the way perceived colour maps onto given patterns of light reflectance may be more invariant than is sometimes thought. After all, it makes evolutionary sense for our perceptual systems to pick out physical invariances when they occur and to translate these into relatively invariant experiences. However even a perfect correlation between perceived qualia and events described by physics would not establish their ontological identity (causation and correlation do not establish identity – see above). Indeed, physical descriptions as such do nothing to explain why one pattern of light reflectances should be perceived as "red," and another as "green," while a pattern of light reflectances in the ultra-violet region is seen as nothing at all (unless one happens to be a bee). Nor do physical descriptions explain the rather arbitrary way the visual system translates electromagnetic energies with wavelengths ordered on a ratio scale into colour categories ordered on a nominal scale. If our experiences simply "mirrored" the world, we would expect the relationships between properties described by physics to be more faithfully preserved in the way such relationships are experienced. To this one must add the many differences in the way given physical properties can be experienced both within and between species (see Velmans, 2000, ch 7 for a review). As van der Heijden et al (1997) note, the view that perceived qualia exist in the world in a way that is free of such biological influences simply does not take the natural sciences seriously.

CONCLUSION

Reductive physicalism rejects first-person evidence, arguing that conscious experiences are nothing more than states of the brain, however they might seem. Having reduced conscious states to brain states, they commonly try to externalise their "qualia", claiming these to be observer-independent, physical properties of the external world. Although reductive physicalism drapes itself in the 'mantle of science,' it routinely ignores the findings of science. For example, it ignores the evidence for the highly specialised nature of human sense organs (sensory physiology), the constructive nature of perception, the complex relationship of experienced qualia to the energies described by physics (psychophysics), the ability of the brain to generate experiences in the absence of the physical energies that those experiences would normally represent (neuropsychology) and the many ways in which human perception differs from that of other animals (comparative psychology). In short, reductive physicalism ignores both the first-person phenomenological evidence regarding the nature of consciousness and the third-person evidence about how it relates to world described by

physics. It is ironic that a philosophy of mind intended to naturalise consciousness gives such an unnatural account of it.

However, it is possible to develop a reflexive model of how consciousness relates to the brain and the physical world that is consistent with both third-person evidence of how the brain works and with first-person evidence of what it is like to have a given experience. Within this model the physical world as perceived (the phenomenal world) is viewed as part of conscious experience not apart from it. While in everyday life we treat this phenomenal world as if it is the "physical world", it is really just one biologically useful representation of what the world is like that may differ in many respects from the world described by physics. How the world as perceived relates to the world as described by physics can be investigated by normal science (e.g. through the study of sensory physiology, psychophysics and so on). While this is an entirely "natural" account of consciousness, it is nonreductive. That is, conscious experiences are really how they seem.

References

- Armstrong DM. A Materialist Theory of Mind. London: Routledge and Kegan Paul. 1968
- Block N. On a confusion about a function of consciousness. Behavioral and Brain Sciences 1995;18(2): 227-272.
- Block N. Biology versus computation in the study of consciousness, Behavioral and Brain Sciences 1997;20(1): 159-166.
- Brugger P. Heautoscopy, epilepsy, and suicide. Journal of Neurology Neurosurgery and Psychiatry 1994;57: 838-839.
- Chalmers D. Facing up to the problem of consciousness. Journal of Consciousness Studies 1995;2(3): 200-219.
- Crick F. The Astonishing Hypothesis: The scientific search for the soul. London: Simon & Schuster. 1994.
- Dennett DC. Consciousness Explained, London: Allen Lane, The Penguin Press. 1991.
- Dennett DC and Kinsbourne M. Time and the observer: The where and when of consciousness in the brain, Behavioral and Brain Sciences 1992;15:183-200.
- James W. Does 'consciousness' exist? Reprinted in: G.N.A. Vesey (ed.) Body and mind: Readings in philosophy, London: Allen & Unwin. 1970 [1904]
- Penfield W and Rassmussen TB. The Cerebral Cortex of Man. Princeton: Princeton University Press. 1950.
- Pribram KH. Languages of the Brain: Experimental paradoxes and principles in neuropsychology. Englewood Cliffs, N.J: Prentice-Hall, 2nd ed. Monterey. Brooks/Cole. 1971.
- Rock I. Indirect Perception. Cambridge, Mass. MIT Press. 1997.
- Searle J. The Rediscovery of the Mind, Cambridge, Mass. MIT Press. 1992.
- Tye M. Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind. Cambridge, Mass. MIT Press. 1995.
- van der Heijden AHC, Hudson PTW and Kurvink AG. On widening the explanatory gap. Behavioral and Brain Sciences 1997;20(1): 157-158.
- Velmans M. Understanding Consciousness. London: Routledge/Psychology Press. 2000.
- Velmans M. Goodbye to reductionism. In: S.Hameroff, A.Kaszniak and A.Scott (eds) Towards a Science of Consciousness II: The Second Tucson Discussions and Debates, Cambridge, Mass. MIT Press. 1998.
- Velmans M. A reflexive science of consciousness, In: Experimental and Theoretical Studies of Consciousness, Ciba Foundation Symposium No.174, Chichester. Wiley. 1993.
- Velmans M. Is human information processing conscious? Behavioral and Brain Sciences 1991;4(4): 651-726.
- Velmans, M. 1990. Consciousness, brain, and the physical world. Philosophical Psychology 3: 77-99.
- Zajonc A. Catching the Light: An Entwined History of Light and Mind. London: Bantam Press.1993.