

Max Velmans

## *How Could Conscious Experiences Affect Brains?*

*In everyday life we take it for granted that we have conscious control of some of our actions and that the part of us that exercises control is the conscious mind. Psychosomatic medicine also assumes that the conscious mind can affect body states, and this is supported by evidence that the use of imagery, hypnosis, bio-feedback and other 'mental interventions' can be therapeutic in a variety of medical conditions. However, there is no accepted theory of mind/body interaction and this has had a detrimental effect on the acceptance of mental causation in science, philosophy and in many areas of clinical practice. Biomedical accounts typically translate the effects of mind into the effects of brain functioning, for example, explaining mind/body interactions in terms of the interconnections and reciprocal control of cortical, neuroendocrine, autonomic and immune systems. While such accounts are instructive, they are implicitly reductionist, and beg the question of how conscious experiences could have bodily effects. On the other hand, non-reductionist accounts have to cope with three problems: (1) The physical world appears causally closed, which would seem to leave no room for conscious intervention. (2) One is not conscious of one's own brain/body processing, so how could there be conscious control of such processing? (3) Conscious experiences appear to come too late to causally affect the processes to which they most obviously relate. This paper suggests a way of understanding mental causation that resolves these problems. It also suggests that 'conscious mental control' needs to be partly understood in terms of the voluntary operations of the preconscious mind, and that this allows an account of biological determinism that is compatible with experienced free will.*

### **What Needs to be Explained**

The assumption that we have a conscious mind that controls our voluntary functions and actions is taken for granted in everyday life and is deeply ingrained in our ethics, politics and legal systems. The potential effect of the mind on the body is also taken for granted in psychosomatic medicine. But *how* the conscious mind exercises its influence is not easy to understand. In principle, there are four

distinct ways in which body/brain and mind/consciousness might enter into causal relationships. There might be physical causes of physical states, physical causes of mental states, mental causes of mental states, and mental causes of physical states. Establishing which forms of causation are effective in *practice* is important, not just for a deeper understanding of mind/body interactions, but also for the proper treatment of some forms of illness and disease.

Within conventional medicine, physical→physical causation is taken for granted. Consequently, the proper treatment for physical disorders is assumed to be some form of physical intervention. Psychiatry takes the efficacy of physical→mental causation for granted, along with the assumption that the proper treatment for psychological disorders may involve psychoactive drugs, neurosurgery and so on. Many forms of psychotherapy take mental→mental causation for granted, and assume that psychological disorders can be alleviated by means of ‘talking cures’, guided imagery, hypnosis and other forms of mental intervention. Psychosomatic medicine assumes that mental→physical causation can be effective (‘psychogenesis’). Consequently, under some circumstances, a physical disorder (for example, hysterical paralysis) may require a mental (psychotherapeutic) intervention. Given the extensive evidence for *all* these causal interactions (cf. readings in Velmans, 1996a), how are we to make sense of them?

### **Clinical Evidence for the Causal Efficacy of Conscious Mental States**

The problems posed by mental→physical causation are particularly acute, as reductionist, materialistic science generally takes it for granted that the operation of physical systems can be entirely explained in physical terms. Yet there is a large body of evidence that states of mind can affect not only subsequent states of the mind but also states of the body. For example, Barber (1984), Sheikh *et al.* (1996), and the readings in Sheikh (2001) review evidence that the use of imagery, hypnosis and biofeedback may be therapeutic in a variety of medical conditions.

Particularly puzzling is the evidence that under certain conditions, a range of *autonomic* body functions including heart rate, blood pressure, vasomotor activity, blood glucose levels, pupil dilation, electrodermal activity and immune system functioning can be influenced by conscious states. In some cases these effects are striking. Baars and McGovern (1996) for example report that,

The global influence of consciousness is dramatized by the remarkable phenomenon of biofeedback training. There is firm evidence that *any* single neurone or *any* population of neurons can come to be voluntarily controlled by giving conscious feedback of their neural firing rates. A small needle electrode in the base of the thumb can tap into a single motor unit — a muscle fibre controlled by one motor neurone coming from the spinal cord, and a sensory fibre going back to it. When the signal from the muscle fibre is amplified and played back as a click through a loud-speaker, the subject can learn to control his or her single motor unit — one among millions — in about ten minutes. Some subjects have learned to play drumrolls on their single motor units after about thirty minutes of practice! However, if the biofeedback signal is not conscious, learning does not occur. Subliminal feedback, distraction from the feedback signal, or feedback via a *habituating* stimulus — all these cases prevent control being acquired. Since this kind of learning only works

for *conscious* biofeedback signals, it suggests again that consciousness creates global access to all parts of the nervous system (p. 75).

The most well accepted evidence for the effect of states of mind on medical outcome is undoubtedly the ‘placebo effect’ — well known to every medical practitioner and researcher. Simply receiving treatment, and having confidence in the therapy or therapist has itself been found to be therapeutic in many clinical situations (cf. Skrabanek and McCormick, 1989; Wall, 1996). As with other instances of apparent mind/body interaction, there are conflicting interpretations of the causal processes involved. For example, Skrabanek and McCormick (1989) claim that placebos can affect illness (how people feel) but not disease (organic disorders). That is, they accept the possibility of mental→mental causation but not of mental→physical causation.

However, Wall (1996) cites evidence that placebo treatments may produce organic changes. Hashish *et al.* (1988) for example, found that use of an impressive ultrasound machine reduced not only pain, but also jaw tightness and swelling after the extraction of wisdom teeth whether or not the machine was set to produce ultrasound. Wall also reviews evidence that placebos can remove the sensation of pain accompanying well-defined organic disorders, and not just the feelings of discomfort, anxiety and so on that may accompany it.

As McMahon and Sheikh (1989) note, the absence of an acceptable theory of mind/body interaction within philosophy and science has had a detrimental effect on the acceptance of mental causation in many areas of clinical theory and practice. Conversely, the extensive evidence for mental causation within some clinical settings forms part of the database that any adequate theory of mind/consciousness–body/brain relationships needs to explain.

### **Some Useful Accounts of Mental Causation**

The theoretical problems posed by mental causation are nicely illustrated by studies of imagery. According to the evidence reviewed by Sheikh *et al.* (1996), imagery can be an effective tool in exercising mental control over ones own bodily states (heart rate, blood pressure, vasomotor activity and so on). It can also affect other states of mind, playing an important role in hypnosis and meditation. But, how could ephemeral images affect the spongy material of brains, and by what mechanism could conscious images affect other conscious states?

In clinical practice, the effects of imagery on brain, body and other conscious experience are often explained to patients in terms of *refocusing and redirection of attention*, linked where plausible to the operation of known biological mechanisms. For example, in their pain control induction programme, Syrjala and Abrams (1996) explain the effectiveness of imagery to patients in terms of the gate-control theory of pain:

Even though the pain message starts in your leg, you won’t feel pain unless your brain gets the pain message. The pain message moves along nerves from where the injury is located to the brain. These nerves enter the spinal cord, where they connect to other nerves, which send information up the spinal cord to the brain. The connections in the spinal cord and brain act like gates. These gates help you to not have to

pay attention to all the messages in your body all the time. For example, right now as you are listening, you do not notice the feelings in your legs, although those feelings are there if you choose to notice them. If you are walking, you might notice feelings in your legs but not in your mouth. One way we block the gates to pain is with medications. Or we can block the gates by filling them with other messages. You do this if you hit your elbow and then rub it hard. The rubbing fills the gate with other messages, and you feel less pain. You've done the same thing if you ever had a headache and you get busy doing something that takes a lot of concentration. You forget about the headache because the gates are full of other messages. Imagery is one way to fill the gate. You can choose to feel the pain if you need to, but any time you like you can fill the gate with certain thoughts and images. Our goal is to find the best gate fillers for you (p. 243).

While this account is nicely judged in terms of its practical value to patients, it does not give much detail about the actual mechanisms involved. Nor does it serve as a general account of mental causation in situations that seem to demand a more sophisticated understanding of the intricate, reciprocal balance of mind/brain/body relationships. The evidence that involuntary processes can sometimes be brought under voluntary control, for example, appears to blur the classical boundary between voluntary and autonomic nervous system functions, and extends the potential scope of top-down processing in the brain. Also, the evidence that imagery can sometimes have bodily effects that resemble the effects of the imaged situations themselves suggest that the conventional, clear distinction between 'psychological reality' and 'physical reality' may not be so clear in the way that these are *responded to* by body and brain. As Kenneth Pelletier (1993) puts it:

Asthmatics sneeze at plastic flowers. People with a terminal illness stay alive until after a significant event, apparently willing themselves to live until a graduation ceremony, a birthday milestone, or a religious holiday. A bout of rage precipitates a sudden, fatal heart attack. Specially trained people can voluntarily control such 'involuntary' bodily functions as the electrical activity of the brain, heart rate, bleeding, and even the body's response to infection. Mind and body are inextricably linked, and their second-by-second interaction exerts a profound influence upon health and illness, life and death. Attitudes, beliefs, and emotional states ranging from love and compassion to fear and anger can trigger chain reactions that affect blood chemistry, heart rate, and the activity of every cell and organ system in the body — from the stomach and gastrointestinal tract to the immune system. All of that is now indisputable fact. However, there is still great debate over the extent to which the mind can influence the body and the precise nature of that linkage. (p. 19)

One productive route to a deeper understanding of such linkages is the traditional biomedical one, involving a fuller understanding of the interconnections and reciprocal control between cortical, neuroendocrine, autonomic and immune systems. These have been extensively investigated within psychoneuroimmunology. Following a detailed review of this research, Watkins (1997) concludes that

It is apparent that the immune system can no longer be thought of as autoregulatory. Virtually every aspect of immune function can be modulated by the autonomic nervous system and centrally produced neuropeptides. These efferent neuro-immuno modulatory pathways are themselves modulated by afferent inputs from

the immune system, the cortex and the limbic emotional centers. Thus the brain and the immune system communicate in a complex bidirectional flow of cytokines, steroids and neuropeptides, sharing information and regulating each other's function. This enables the two systems to respond in an integrated manner to environmental challenges, be they immunological or behavioral, and thereby maintain homeostatic balance (p. 15).

### So Why Does Mental Causation Remain a Problem?

Such innovative findings and their practical consequences for the development of 'mind-body medicine' demand careful investigation. It is important to note however that such explanatory accounts routinely translate *mind*-body interactions into *brain*-body interactions. Unless one is prepared to accept that mind and consciousness are *nothing more* than brain processes<sup>1</sup> this finesses the classical mind/body problems that are *already* posed by *normal* voluntary, 'mental' control. How imagery might affect autonomic or immune system functioning is mysterious, but how a conscious wish to lift a finger makes that finger move is equally mysterious. Why? There are many reasons, but I will focus on just three.

#### 1. *The physical world appears causally closed*

As noted above, it is widely accepted in science that the operation of physical systems can be entirely explained in physical terms. For example, if one examines the human brain from an external third-person perspective one can, in principle, trace the effects of input stimuli on the central nervous system all the way from input to output, without finding any 'gaps' in the chain of causation that consciousness might fill. Indeed, the *neural correlates* of consciousness would fill any 'gaps' that might potentially be filled by consciousness in the activities of brain. In any case, if one inspects the operation of the brain from the outside, no subjective experience can be observed at work. Nor does one need to appeal to the existence of subjective experience to account for the neural activity that one *can* observe. The same is true if one thinks of the brain as a functioning system described in information processing terms rather than neural terms. Once the processing within a system required to perform a given function is sufficiently well specified in procedural terms, one does not have to add an 'inner conscious life' to make the system work. In principle, the same function, operating to the same specification, could be performed by a non-conscious machine.<sup>2</sup>

- 
- [1] Although variants of eliminative/reductive physicalism and functionalism (that consciousness is nothing more than a state or function of the brain) are commonly adopted in current philosophy and science, the reduction of conscious phenomenology to brain states or functions faces well-recognized difficulties. I present a detailed analysis of the strengths and weaknesses of various eliminative, reductive and emergent forms of physicalism, along with psychofunctionalism (functionalism in cognitive psychology) and computational functionalism (functionalism in philosophy and AI) in Velmans (2000) chs. 3, 4 and 5. On-line papers addressing many of the difficulties, for example in the work of Searle, Dennett, Armstrong, Block and Tye are also available from the CogPrints archive (<http://cogprints.soton.ac.uk/>) — see Velmans (1998; 2001a,b). Given the current prevalence of physicalism I also summarize some of my reasons for not adopting it in the Appendix below.
- [2] Note that being physically closed does not preclude 'downward causation'. Higher order brain states or functions may for example constrain lower order brain states and functions, for example in the way

2. *One is not conscious of one's own brain/body processing, so how could there be conscious control of such processing?*

How 'conscious' is conscious, voluntary control? It is surprising how few people bother to ask.<sup>3</sup> One might be aware of the fact *that* relaxing imagery can lower heart rate, but one has no awareness of *how* it does so, nor, in biofeedback, does one have any awareness of how consciousness might control the firing of a single motor neurone. One isn't even conscious of *how* to control the articulatory system in everyday 'conscious speech'! Speech production is one of the most complex tasks humans are able to perform. Yet, one has no awareness whatsoever of the motor commands issued from the central nervous system that travel down efferent fibres to innervate the muscles, nor of the complex motor programming that enables muscular co-ordination and control. In speech, for example, the tongue may make as many as twelve adjustments of shape per second — adjustments which need to be precisely coordinated with other rapid, dynamic changes within the articulatory system. According to Lenneberg (1967), within one minute of discourse as many as ten to fifteen thousand neuromuscular events occur. Yet only the *results* of this activity (the overt speech) normally enters consciousness.

Preconscious speech control might of course be the result of *prior* conscious activity, for example, planning *what* to say might be conscious, particularly if one is expressing some new idea, or expressing some old idea in a novel way. Speech production is commonly thought to involve hierarchically arranged, semantic, syntactic and motor control systems in which communicative intentions are translated into overt speech in a largely top-down fashion. Planning *what* to say and translating nonverbal conceptual content into linguistic forms requires effort. But to what extent is such planning conscious? Let us see.

A number of theorists have observed that periods of conceptual, semantic and syntactic planning are characterized by gaps in the otherwise relatively continuous stream of speech (Goldman-Eisler, 1968; Boomer, 1970). The neurologist John Hughlings Jackson, for example, suggested that the amount of planning required depends on whether the speech is 'new' speech or 'old' speech. Old speech (well-known phrases, etc.) requires little planning and is relatively continuous. New speech (saying things in a new way) requires planning and is characterized by hesitation pauses. Fodor *et al.* (1974) point out that breathing pauses also occur (gaps in the speech stream caused by the intake of breath). However, breathing pauses do not generally coincide with hesitation pauses.

Breathing pauses nearly always occur at the beginnings and ends of major linguistic constituents (such as clauses and sentences). So these appear to be coordinated with the syntactic organization of such constituents into a clausal or sentential structure. Such organization is largely automatic and preconscious. By contrast, hesitation pauses tend to occur within clauses and sentences and appear

---

that computer software constrains and controls the switching in the hardware of the machine. The software, like the higher order functioning of the brain is best described in functional terms (e.g. as an information processing system), but this does not alter the fact that the software is entirely embodied in the physical hardware, and exercises its causal effects through its embodiment in that hardware.

[3] See the initial discussion of this issue in Velmans (1991a).



to be associated with the formulation of ideas, deciding which words best express one's meaning, and so on. If this analysis is correct, conscious planning of *what* to say should be evident during hesitation pauses — and a little examination of what one experiences during a hesitation pause should settle the matter. Try it. During a hesitation pause one might experience a certain sense of effort (perhaps the effort to put something in an appropriate way). But nothing is revealed of the *processes* that formulate ideas, translate these into a form suitable for expression in language, search for and retrieve words from memory, or assess which words are most appropriate. In short, no more is revealed of conceptual or semantic planning in hesitation pauses than is revealed of syntactic planning in breathing pauses. The fact that a process demands processing *effort* does not ensure that it is *conscious*. Indeed, there is a sense in which one is only conscious of what one wants to say *after one has said it!*

It is particularly surprising that the same may be said of *conscious verbal thoughts*. That is, the same situation applies if one formulates one's thoughts into 'covert speech' through the use of phonemic imagery, prior to its overt expression. Once one *has* a conscious verbal thought, manifested in experience in the form of phonemic imagery, the complex cognitive processes required to generate that thought, including the processing required to encode it into phonemic imagery *have already operated*. In short, covert speech and overt speech have a similar relation to the planning processes that produce them. In neither case are the complex antecedent processes available to introspection. It should be clear that this applies equally to the processes that generate the detailed spatial arrangement, colours, shapes, sizes, movements and accompanying sounds and smells of an imaged visual scene.

### *3. Conscious experiences appear to come too late to causally affect the processes to which they most obviously relate*

In the production of overt speech and covert speech (verbal thoughts) the conscious experience that we normally associate with such processing *follows* the processing to which it relates. Given this, in what *sense* are these 'conscious processes' conscious? The same question can be asked of that most basic of conscious voluntary processes, *conscious volition itself*.

It has been known for some time that voluntary acts are preceded by a slow negative shift in electrical potential (recorded at the scalp) known as the 'readiness potential', and that this shift can precede the act by up to one second or more (Kornhuber and Deeke, 1965). In itself, this says nothing about the relation of the readiness potential to the *experienced wish* to perform an act. To address this, Libet (1985) asked subjects to note the instant they experienced a wish to perform a specified act (a simple flexion of the wrist or fingers) by relating the onset of the experienced wish to the spatial position of a revolving spot on a cathode ray oscilloscope, which swept the periphery of the face like the sweep-second hand of a clock. Recorded in this way, the readiness potential preceded the voluntary act by around 550 milliseconds, and preceded the experienced wish (to flex the wrist or fingers) by around 350 milliseconds (for spontaneous acts involving

no preplanning). This suggests that, like the act itself, the experienced wish (to flex one's wrist) may be one output from the (prior) cerebral processes that actually select a given response. If so, 'conscious volition' may be no more necessary for such a (preconscious) choice than the consciousness of one's own speech is necessary for its production.<sup>4</sup> The same is likely to apply to more complex voluntary acts, such as the voluntary control of autonomic functions through imagery and biofeedback discussed above.<sup>5</sup>

### The Current Theoretical Impasse

As noted, there is extensive experimental and clinical evidence that conscious experiences can affect brain/body processes, and the importance of conscious experience is rightly taken for granted in everyday life. In one sense this can be explained by a more sophisticated biomedical understanding of mind/brain/body relationships. But in a deeper sense, current attempts to understand the role of conscious experience face an impasse. How can experiences have a causal influence on a physical world that is causally closed? How can one consciously control something that one is not conscious of? How can experiences affect processes that *precede* them? Dualist-interactionist accounts of the consciousness–brain relationship, in which an autonomously existing consciousness influences the brain, do not even recognize these 'how' problems let alone address them. Materialist reductionists attempt to finesse such problems by challenging the accuracy, causal efficacy and even the existence of conscious experiences. This evades the need to address the 'how' questions, but denies the validity of the clinical evidence and defies common sense. I have given a detailed critique of the many variants of dualism and reductionism elsewhere and will not repeat it here.<sup>6</sup> In what follows I suggest a way through the impasse that is neither dualist nor reductionist.<sup>7</sup>

### Ontological Monism Combined with Epistemological Dualism

How can one reconcile the evidence that conscious experiences are causally effective with the principle that the physical world is causally closed? One

- 
- [4] As Libet observed, the experienced wish *follows* the readiness potential, but *precedes* the motor act itself (by around 200 msec) — time enough to consciously *veto* the wish before executing the act. In a manner reminiscent of the interplay between the libidinous desires arising from Freud's unconscious *id* and the control exercised by the conscious *ego*, Libet suggested that the *initiation* of voluntary act and the accompanying wish are developed preconsciously, but consciousness can then act as a form of censor which decides whether or not to carry out the act. While this is an interesting possibility, it does invite an obvious question. If the wish to perform an act is developed preconsciously, why doesn't the decision to censor the act have its own preconscious antecedents? Libet (1996) argues that it *might* not need to do so as voluntary control imposes a change on a wish that is already conscious. Yet, it seems very odd that a wish to do something has preconscious antecedents while a wish not to do something does not. As it happens, there is evidence that bears directly on this issue. Karrer *et al.* (1978), and Konttinen and Lyytinen (1993), for example, found that *refraining* from irrelevant movements is associated with a slow *positive-going* readiness potential.
- [5] This could be tested using Libet's procedures, by examining the relation of the readiness potential to an experienced wish to control a given bodily function via imagery or biofeedback.
- [6] See Velmans (2000) chs. 2, 3, 4 and 5 and the Appendix below.
- [7] In the space available I can give only an introduction to how one might resolve these problems. A more detailed treatment is given in Velmans (2000) ch. 11.



simple way is to accept that for each individual there is *one* ‘mental life’ but *two* ways of knowing it: first-person knowledge and third-person knowledge. From a first-person perspective conscious experiences appear causally effective. From a third-person perspective the same causal sequences can be explained in neural terms. It is not the case that the view from one perspective is right and the other wrong. These perspectives are complementary. The differences between how things appear from a first- versus a third-person perspective has to do with differences in the *observational arrangements* (the means by which a subject and an external observer access the subject’s mental processes).

Let’s see how this might work in practice. Suppose you have a calming image of lying in a green field on a summer’s day, and you can feel the difference this makes in producing a relaxed state, slowing your breathing, removing the tension in your body and so on. You give a causal account of what is going on, based on what you experience. From my external observer’s perspective, I can also observe what is going on — but what I observe is a little different. I can measure the effects on your breathing and muscle tension, but no matter how closely I inspect your brain I cannot observe your experienced image. The closest I can get to it are its neural correlates in the visual system, association areas and so on.<sup>8</sup> Nevertheless, if I could observe all the neurophysiological events operating in your brain to produce your relaxed bodily state, I could give a complete, physical account of what is going on. So, now you have a first-person account of what is going on that makes sense to you and I have a third-person account of what is going on that makes sense to me. How do these relate? To understand this we need to examine the relation of your visual image to its neural correlates with care.

### *The neural correlates of conscious experience*

Although we know little about the physical nature of the neural correlates of conscious experiences, there are three plausible, functional constraints imposed by the phenomenology of consciousness itself. Normal human conscious experiences are representational (phenomenal consciousness is always *of something*).<sup>9</sup>

- 
- [8] The neural correlates of a given experience accompany or *co-occur* with given experiences, and are by definition as close as one get to those experiences from an external observer’s perspective. This differentiates them from the antecedent causes (such as the operation of selective attention, binding, etc.) which may be thought of as the necessary and sufficient *prior* conditions for given experiences in the human brain.
- [9] My assumption that normal conscious experiences are representational is driven by a Critical Realist epistemology (developed in Velmans, 2000, ch. 7) and not by any commitment to the view that mental states are nothing more than computations on representations (a thesis that is currently in dispute). While I do not have space to develop the case for Critical Realism here, it is worth noting that there is nothing mysterious about experiences being representations of entities and events outside or within our bodies and brains that differ in some respects from the alternative representations of those entities and events given by science (e.g. by physics). Perceptual processes are likely to have developed in response to evolutionary pressures, and select, attend to and interpret information in accordance with human adaptive needs. Consequently, they only need to model a subset of the available information. At the same time our perceptual models must be useful, otherwise it is unlikely that human beings would have survived. Given this, it seems reasonable to assume that, barring illusions or hallucinations, the experiences produced by perceptual processing are partial, approximate but nonetheless useful representations of what is ‘really there’. [*footnote continued over page*]

Given this, it is reasonable to assume that the neural correlates of such experiences are also representational states.

Although this assumption has not always been made explicit in theories of consciousness it is largely taken for granted in psychological theory. Psychophysics, for example, takes it for granted that for any discriminable aspect of experience (a just noticeable change in brightness, colour, pitch and so on) there will be a correlated change in some state of the brain. It follows from this that the information encoded in experience (in terms of discriminable differences) will also be encoded in the brain. The same is true for the more complex contents of consciousness, in the many cognitive theories that associate (or identify) such contents with information stored in primary (working) memory, information at the focus of attention, information in a global workspace and so on.

A representational state must, of course, represent *something*, that is it must have a given content. For a given physical state to be the correlate of a given experience it is plausible to assume that it represents the *same* thing (otherwise it would not be the correlate of *that* experience).

Finally, for a physical state to be the correlate of a given experience, it is reasonable to suppose that it has the same 'grain'. That is, for every discriminable attribute of experience there will be a distinct, correlated, physical state. As each experience and its physical correlate represents the same thing it follows that each experience and its physical correlate encodes the same information about that thing. That is, they are representations with the same *information structure*.<sup>10, 11</sup>

---

The view that some conscious experiences are representational in the sense of being 'intentional' (that they are *of* something) has in any case been widely accepted in philosophy of mind since Brentano reintroduced this medieval notion in the nineteenth century. According to some philosophers not all conscious experiences are intentional. Searle (1994b) for example maintains that 'a feeling of pain or a sudden sense of anxiety, where there is no object of the anxiety, are not intentional' (p. 380). In Velmans (1990; 2000) I argue that a conscious experience does not have to be about a specific external object for it to be representational. It may for example represent a state of one's own body or it may represent a global *reaction* to a real, imagined or remembered event. A feeling of pain, for example, represents (in one's first-person experience) actual or potential damage to the body, and it is usually quite accurate in that it is normally subjectively located at or near the site of body damage. A feeling of anxiety is a first-person representation of a state of one's own body and brain that signals actual or potential danger, and so on. Viewed this way, *all* conscious states are about something. On this issue, I adopt the same stance as that developed by Tye (1995).

[10] This assumption of conscious experience/neural correlate functional equivalence (defined in information processing terms) is a point of convergence between otherwise widely divergent theories (physicalism, functionalism, dual-aspect theory). As Gardner (1987) points out, the assumption that mental processes operate on representations lies at the foundation of cognitive science. However, the claim that the neural correlates of conscious states are representations begs no questions about the forms that these representations might take, or about how mental processes operate on them. Representations might be iconic, propositional, feature sets, prototypes, procedural, localized, distributed, static or dynamic, or whatever. Operations might be formal and computational, or more like the patterns of shifting weights and probabilities that determine the activation patterns in neural networks. I suggest that the correlates of consciousness represent what the phenomenology itself represents, irrespective of how the correlates embody those representations.

[11] This approach has its origins in Spinoza's dual-aspect theory, which I developed into a naturalized, dual-aspect theory of information in Velmans (1991a,b; 1993; 1996; 2000). This dual-aspect theory of information also has similarities to that adopted by Chalmers (1996) (see Velmans, 2000, p. 281, note 5, for a summary of both the similarities and differences). Note that having an identical referent

If these assumptions are well founded, your experience and the neural correlates that I observe will relate to each other in a very precise way. What you experience takes the form of visual or other imagery accompanied by feelings about lying on the grass on a summery day. What I observe *is the same information* (about the visual scene) encoded in the physical correlates of what you experience in your brain. The information structure of what you and I observe is identical, although it is displayed or ‘formatted’ in very different ways. From your point of view, the only information you have about your own state of mind is the imagery and accompanying feelings that you experience. From my point of view, the only information you have (about your own state of mind) is the information I can see encoded in your brain. The way your information (about your own state) is displayed appears to be very different to you and me for the reason that the ‘observational arrangements’ by which we access that information are entirely different. From my external, third-person perspective I can only access the information encoded in your mind/brain by means of my visual or other exteroceptive systems aided by appropriate equipment. With these means I can detect the information displayed in the form of neural encodings, but not in the form of accompanying experiences. While you maintain your focus on the imaged scene, you cannot observe its neural correlates in your own brain (you would need to use my equipment for that). Nevertheless, the information in those correlates displays ‘naturally’,<sup>12</sup> in the form of the imaged scene that you experience.

But what is your mind *really* like? From my ‘external observer’s perspective’, can I assume that what you experience is really nothing more than the physical correlates that I can observe? From my external perspective, do I know what is going on in your mind/brain/consciousness better than you do? No. I know something about your mental states that you do not know (their physical embodiment). But you know something about them that I do not know (their manifestation in your experience). Such first- and third-person information is *complementary*. We need your first-person story and my third-person story for a complete account of what is going on. If so, the nature of the mind is revealed as much by how it appears from one perspective as the other. It is not *either* physical *or* conscious experience, it is at once physical *and* conscious experience (depending on the observational arrangements). For lack of a better term we may describe this nature as *psychophysical*.<sup>13, 14</sup> If we combine this with the representational features above,

and information structure does not mean that experiences are *nothing more than* their neural correlates (as eliminativists and reductionists assume). A filmed version of the play *Hamlet*, recorded on videotape, for example, may have the same sequential information structure as the same play displayed in the form of successive, moving pictures on a TV screen. But it is obvious that the information on the videotape is not ontologically identical to the information displayed on the screen. In this instance, the same information is embodied in two different ways (patterns of magnetic variation on tape versus patterns of brightness and hue in individual pixels on screen) and it is displayed or ‘formatted’ in two different ways (only the latter display is in visible form).

[12] I assume it to be a *natural* fact about the world that certain forms of neural activity are accompanied by conscious experiences. Consequently, when such neural activities (the correlates) occur in one’s brain one has the corresponding experiences. I also assume that the formatting of neurally encoded

we can say that mind is a psychophysical process that encodes information, developing over time.

### **An Initial Way to Make Sense of the Causal Interactions between Consciousness and Brain**

This brief analysis of how first- and third-person accounts relate to each other can be used to make sense of the different *forms* of causal interaction that are taken for granted in everyday life or suggested in the clinical and scientific literature. Physical→physical causal accounts describe events from an entirely third-person perspective (they are ‘pure third-person accounts’). Mental→mental causal accounts describe events entirely from a first-person perspective (they are ‘pure first-person accounts’). Physical→mental and mental→physical causal accounts are mixed-perspective accounts employing perspectival switching

---

information relates to the formatting of corresponding, phenomenally encoded information in an orderly way, with discoverable neural-state space/phenomenal space mappings. An obvious example would be the way that information about spatial location and extension encoded in the brain is mapped into the 3D phenomenal space that we ordinarily experience. In vision, some progress has already been made in the discovery of such mappings (see the Special Issue on the work of Roger Shepard in *Behavioural and Brain Sciences*, 24 (4), 2001). While neural state/phenomenal state mappings are likely to differ in different sense modalities (e.g. vision versus audition) and even between different features of a given modality (e.g. colour versus spatial location and extension) there may also be shared, underlying principles (cf. Stoffregen and Benoit, 2001).

- [13] The struggle to find a model or even a form of words that somehow captures the dual-aspect nature of mind is reminiscent for example of wave-particle complementarity in quantum mechanics — although this analogy is far from exact. Light either appears to behave as electromagnetic waves or as photon particles depending on the observation arrangements. It does not make sense to claim that electromagnetic waves really *are* particles (or vice versa). A complete understanding of light requires both complementary descriptions — with consequent struggles to find an appropriate way of characterizing the nature of light and other QM phenomena which encompass both descriptions (‘wave-packets’, ‘electron clouds’ and so on). This has not prevented physics from developing very precise accounts of light viewed *either* as waves *or* as particles, together with precise formulae for relating wave-like properties (such as electromagnetic frequency) to particle-like ones (such as photon energy). If first- and third-person accounts of consciousness and its physical correlates are complementary and mutually irreducible, an analogous ‘psychological complementarity principle’ might be required to understand the nature of mind. A more detailed discussion of how psychological complementarity relates to physical complementarity is given in Velmans (2000) ch. 11, note 19.
- [14] At the macrocosmic level, the relation of electricity to magnetism also provides a clear parallel to the form of dual-aspect theory I have in mind. If one moves a wire through a magnetic field this produces an electrical current in the wire. Conversely, if one passes an electrical current through a wire this produces a surrounding magnetic field. But it does not make sense to suggest that the current in the wire is nothing more than the surrounding magnetic field, or vice-versa (reductionism). Nor is it accurate to suggest that electricity and magnetism are energies of entirely different kinds that happen to interact (dualist-interactionism). Rather these are two manifestations (or ‘dual-aspects’) of *electromagnetism*, a more fundamental energy that grounds and unifies both, described with elegance by Maxwell’s Laws. Analogously, phenomenally encoded information and its correlated neurally encoded information may be two manifestations (or ‘dual-aspects’) of a more fundamental, psychophysical mind, and their relationship may, in time, be describable by neurophenomenological laws (see also note 12, above). It goes without saying that a *fully* satisfying psychophysical account of any given mental state would have to specify how given complementary first- and third-person descriptions relate to each other with precision (perhaps with the elegance of Maxwell’s Laws). However, such empirical relationships can only be discovered by neuropsychological research, and for the present I am only concerned with the *form* that causal accounts based on such research might need to take to resolve this aspect of the ‘causal paradox’.

(Velmans, 1996b). Such accounts start with a description of causes viewed from one perspective (either first- or third-person) and then switch to a description of effects viewed from the other perspective. To understand such accounts one first has to acknowledge that a perspectival switch has taken place.

Physical→mental causal accounts start with events viewed from a third-person perspective and switch to how things appear from a first-person perspective. For example, a causal account of visual perception starts with a third-person description of the physical stimulus and the visual system but then switches to a first-person account of what the subject experiences. Mental→physical causal accounts switch the other way. From your subjective point of view, for example, the imagery that you experience is causing your heart rate to slow down and your body to relax (effects that I can measure). If I could identify the exact neural correlates of what you experience, it might be possible for me to give an entirely third-person account of this sequence of events (in terms of higher order neural representations having top-down effects on other brain and body states). But the mixed-perspective account actually gives you a more immediately useful description of what is going on in terms of the things that you can do (maintain that state of mind, deepen it, alter it, and so on).

In principle, complementary first- and third-person sources of information can be found whenever body or mind/brain states are represented in some way in subjective experience. A patient might for example have insight into the nature of a psychological problem (via feelings and thoughts), that a clinician might investigate by observing his/her brain or behaviour. In medical diagnosis, a patient might have access to some malfunction via interoceptors, producing symptoms such as pain and discomfort, whereas a doctor might be able to identify the cause via his/her exteroceptors (eyes, ears and so on) supplemented by medical instrumentation. As with conscious states and their neural correlates the clinician has access to the physical embodiment of such conditions, while the patient has access to how such conditions are experienced. In these situations, neither the third-person information available to the clinician nor the first-person information available to the patient is *automatically* privileged or 'objective' in the sense of being 'observer-free'. The clinician merely reports what he/she observes or infers about what is going on (using available means) and the patient does likewise. Such first- and third-person accounts of the subject's mental life or body states are complementary, and mutually irreducible. *Taken together*, they provide a global, psychophysical picture of the condition under scrutiny.

### **Conscious Experiences are Current, Global Representations Formed by the Mind/Brain**

The above, I hope, gives an initial indication of how one can reconcile the evidence that conscious experiences appear causally effective with the principle that the physical world is causally closed. But there are two further, equally perplexing problems. (1) How can conscious experiences be causally effective if they come too late to affect the mind/brain processes to which they most obviously relate? (2) How can the contents of consciousness affect brain and body



states when one is not conscious of the biological processes that govern those states?

I suggest that to make sense of these puzzles, one has to begin by accepting the facts rather than sweeping them under some obscuring theoretical carpet. Why do experiences come too late to affect the mind/brain processes to which they most closely relate? For the simple reason that experiences relate most closely to the processes that *produce* them. Visual perception becomes ‘conscious’ once visual processing results in a conscious visual experience, cognitive processing becomes ‘conscious’ once it produces the inner speech that forms a conscious thought and so on. Once such experiences arise the processes that have produced them have already taken place. Given this, what is consciousness actually contributing to conscious perception, to conscious speech, to conscious thought, to conscious voluntary control, and so on?<sup>15</sup>

As noted above, I am proceeding on the assumption that conscious experiences are representations. Some experiences represent states of the external world (exteroceptive experiences), some represent states of the body (interoceptive experiences), and some represent states of the mind/brain itself (volitions, thoughts about thoughts, etc.). Experiences can also represent past, future, real and imaginary events, for example in the form of thoughts and images.

Whatever their representational content, current experiences also tell one something important about the current state of one’s own mind/brain — that it currently has percepts, feelings, thoughts, images, etc., of a given type, and that it has formed current representations with that particular content, as opposed to any

---

[15] In Velmans (1991) I argue that there are three distinct senses in which a process may be said to be conscious. It can be conscious (a) in the sense that one is conscious *of* it, (b) in the sense that it *results* in a conscious experience, and (c) in the sense that consciousness *causally affects* that process. We do not have introspective access to how the preconscious cognitive processes that enable thinking produce individual, conscious thoughts in the form of ‘inner speech’. However, the content of such thoughts and the sequence in which they appear does give some insight into the way the cognitive processes (of which they are manifestations) operate over time in problem solving, thinking, planning and so on. Consequently such cognitive processes are partly conscious in sense (a), but only insofar as their detailed operation is made explicit in conscious thoughts, thereby becoming accessible to introspection. Many psychological processes are conscious in sense (b), but not in sense (a) — that is, we are not conscious of how they operate, but we are conscious of their *results*. This applies to perception in all sense modalities. When consciously reading this sentence, for example, you become aware of the printed text on the page, accompanied, perhaps, by inner speech (phonemic imagery) and a feeling of understanding (or not), but you have no introspective access to the processes which enable you to read. Nor does one have introspective access to the *details* of most other forms of cognitive functioning, for example to the detailed operations which enable ‘conscious’ learning, remembering, engaging in conversations with others and so on.

Crucially, having an experience that gives some introspective access to a given process, or having the results of that process manifest in an experience, says nothing about whether that experience *carries out* that process. Whether a process is ‘conscious’ in sense (a) or (b) needs to distinguished from whether it is conscious in sense (c). Indeed, it is not easy to envisage how the experience that makes a process conscious in sense (a) or (b), *could* make it conscious in sense (c). Consciousness *of* a physical process does not make consciousness responsible for the operation of that process (watching a kettle does not determine when it comes to the boil). So, how could consciousness *of* a mental process carry out the functions of that process? Alternatively, if conscious experience *results* from a mental process it arrives *too late* to carry out the functions of that process (see Velmans, 2000, ch. 9 for a more detailed discussion).



others. For example, the thoughts that enter consciousness at a given moment ‘represent’ the current state of one’s own cognitive system in that they reveal *which* of many possible cognitions are currently at the focus of attention in a reportable form. If your thoughts are conscious, and I ask you what you are thinking about, you can tell me. Likewise, your visually imaged peaceful world and your conscious feelings about it represent a current, voluntarily produced representational state (and affective responses to it) within your own visual, cognitive and affective systems — and if I want to know what that is like, you can tell me.

Why don’t we have more detailed experiences of the processes which produce such conscious experiences, or of the detailed workings of our own bodies, minds and brains? Because for normal purposes we don’t need them! Our primary need is to interact successfully with the external world and with each other — and for that, the processes by which we arrive at representations of ourselves in the world, or which govern the many internal, adaptive adjustments we have to make are best left on ‘automatic’. This is exemplified by the well-accepted transition of skills from being conscious to being nonconscious as they become well learnt (as in reading or driving a car). The global representations that we have of ourselves in the world nevertheless provide a useful, reasonably accurate representation of what is going on.<sup>16</sup>

### **How to Make Sense of the Causal Role of the Contents of Consciousness**

As noted above, normal experiences are *of* something, i.e. they represent entities, events and processes in the external world, the body and the mind/brain itself. In everyday life we also behave as ‘naïve realists’. That is we take the events we experience to *be* the events that are actually taking place, although sciences such as physics, biology and psychology might represent the same events in very different ways. For everyday purposes, the assumption that the world just *is* as we experience it to be serves us well. When playing billiards, for example, it is safe to assume that the balls are smooth, spherical, coloured, and cause each other to move by mechanical impact. One only has to judge the precise angle at which the white ball hits the red ball to pocket the red. A quantum mechanical description of the microstructure of the balls or of the forces they exert on each other won’t improve one’s game.

That said, the experienced world is not the world *in itself* — and it is not our experience *of* the balls that governs the movement of the balls themselves. Balls as-experienced and their perceived interactions are global *representations* of autonomously existing entities and their interactions, and conscious representations of such entities or events can only be formed once they exist, or after they have taken place. The same may be said of the events and processes that we experience to occur in our own bodies or minds/brains. When we withdraw a

---

[16] It is reasonable to suppose that the detail of conscious representation has been tailored by evolutionary pressures to be useful for everyday human activities (although these remain global, approximate and species-specific). To obtain a more intricate knowledge of the external world, body or mind/brain we usually need the assistance of scientific instruments. A much fuller analysis of these points is given in Velmans (2000), ch. 7.

hand quickly from a hot iron we experience the pain (in the hand) to cause what we do, but the reflex action actually takes place before the experience of pain has time to form. This can also happen with voluntary movements. Suppose, for example, that you are required to press a button as soon as you feel a tactile stimulus applied to your skin. A typical reaction time is 100 ms or so. It takes only a few milliseconds for the skin stimulus to reach the cortical surface, but Libet, *et al.* (1979) found that awareness of the stimulus takes at least 200 ms to develop. If so, the reaction must take place preconsciously, although we *experience* ourselves as responding *after* we feel something touching the skin. The mind/brain requires time to form a conscious representation of a pain or of something touching the skin and of the subsequent response. Although the conscious representations accurately place the cause (the stimulus) before the effect (the response), once the representations are formed, both the stimulus and the response have already taken place.<sup>17</sup>

Just as the interactions amongst experienced billiard balls represent causal sequences in the external world, but are not the events themselves, experienced interactions between our sensations, thoughts, images and actions represent causal sequences within our bodies and brains, but are not the events themselves. The thoughts, images and feelings that appear in our awareness are both *generated by* processes in our bodies and mind/brains and *represent* the current states of those processes. Thoughts and images represent the ongoing state of play of our cognitive systems; feelings represent our internal (positive and negative) reactions to and judgments about events (see Mangan, 1993, and the discussion above).

In sum, conscious representations of inner, body and external events are not the events themselves, but they generally represent those events and their causal interactions sufficiently well to allow a fairly accurate understanding of what is happening in our lives. Although they are only *representations* of events and their causal interactions, for everyday purposes we can take them to *be* those events and their causal interactions. When we play billiards we can line up a shot without the assistance of physics. Although our knowledge of our own inner states is not incorrigible, when we experience our verbal thoughts expressed in covert or overt speech, we usually know all we need to know about what we currently think — without the assistance of cognitive psychology. When we experience ourselves to have acted out of love or fear, we usually have an adequate understanding of our motivation — although a neuropsychologist might find it

[17] Although conscious experiences arise too late to play a causal role in the processes with which they are most closely associated (those that produce them), once they arise, they are not, of course, too late to play a causal role in other, *subsequent* mind/brain/body states or activities. A pain in the tooth, for example, might persist long enough to force one to the dentist. A desire for employment might lead one to make a job application and so on. However such forms of mental<sup>1</sup>physical causation still face the problem (already discussed) that the physical world is causally closed. For example, the physical movements that take one to the dentist can be explained by the way that the neural correlates of the pain enter into the control of motor systems, the desire for employment in terms of a goal state that is represented in one's CNS and so on. Such forms of mental causation can, however, be understood as 'mixed-perspective' causal accounts of the kind described above. See also the extensive treatment of this particular issue in discussion with Rakover (1996) in Velmans (1996b).

useful to give a third-person account of emotion in terms of its neural substrates in the neocortical, subcortical, diencephalic, midbrain and pontine-medullary brainstem systems (Watt, 2000). When we image ourselves in green grass on a summer's day and feel relaxed we are usually right to assume that the mental state that is represented in our imagery has produced a real bodily effect. For everyday life, it doesn't matter that we don't understand how such imaged scenarios are constructed by preconscious mental processes or exercise top-down control in the mind/brain/body system. It is not the case that a lower level (microscopic) representation is always better than a macroscopic one (in the case of billiard balls). Nor are third-person accounts always better than first-person ones (in describing or attempting to control our thoughts, images and emotions). The value of a given representation, description or explanation can only be assessed in the light of the purposes for which it is to be used.

### Who's in Control?

The difference between voluntary and involuntary bodily functions is accepted wisdom, enshrined in the voluntary/autonomic nervous system distinction in medical texts. As we have seen above, some processes that are normally involuntary can also become partly voluntary once they are represented in consciousness (via biofeedback, imagery and so on). But if we don't have a detailed conscious awareness of the workings of our own bodies and brains and if consciousness comes too late to affect the processes to which it most closely relates how can this be? Consider again the dilemma posed by Libet *et al*'s (1979) experiments on the role of conscious volition described above. If the brain prepares to carry out a given action around 350 milliseconds before the conscious wish to act appears, then how could that action be 'conscious' and how could it be 'voluntary'? Doesn't the preceding readiness potential indicate that the action is determined preconsciously and automatically by processing in the mind/brain?

Let us consider the 'conscious' aspect first. The decision to act (indexed by the readiness potential) is taken preconsciously but it becomes conscious at the moment that it manifests *as* a wish to do something in conscious experience. The wish then becomes conscious in the same way that your perception of this WORD is conscious. Like the wish, once you become conscious of this WORD, the physical, syntactic and semantic analyses required to recognize it have already taken place. Nonetheless, once you become conscious of the wish or the WORD the mental/brain processes make a transition from a preconscious to a conscious state — and it is only when this happens that you consciously realize what is going on.<sup>18</sup>

But how could an act that is executed *preconsciously* be 'voluntary'? Voluntary actions imply the possibility of choice, albeit choice based on available external and internal information, current needs and goals. Voluntary actions are also potentially flexible and capable of being novel. In the psychological

[18] I do not have space to develop this theme in more detail here. In Velmans (2000), chs. 10, 11 and 12, I develop a broader 'reflexive monist' philosophy in which the function of consciousness is to 'real-ize' the world. That is, once an entity, event or process enters consciousness it becomes *subjectively real*.

literature these properties are traditionally associated with controlled rather than automatic processing or with focal-attentive rather than pre-attentive or non-attended processing.<sup>19</sup> Unlike automatic or pre-attentive processing, both controlled processing (in the execution of acts) and focal-attentive processing (in the analysis of input) are thought to be ‘conscious’. None of the above argues against such traditional wisdom. In Libet’s experiments the conscious experience appears around 350 milliseconds after the onset of preconscious processes that are indexed by the readiness potential. This says something about the timing of the conscious experience in relation to the processes that generate it and about its restricted role once it appears. However, it does not argue against the voluntary nature of that preconscious processing. On the contrary, the fact that the act consciously feels as if it is voluntary and controlled suggests that the processes which have generated that experience *are* voluntary and controlled, as conscious experiences generally provide reasonably accurate representations of what is going on (see above). This applies equally to the voluntary nature of more complex, mental processing such as the self-regulating, self-modifying operations of our own psychophysical minds evidenced by the effects of conscious imagery, meditation and biofeedback. In short, I suggest that the feeling that we are free to choose or to exercise control is compatible with the nature of what is actually taking place in our own central nervous system, following processes that select amongst available options, in accordance with current needs, goals, available strategies, calculations of likely consequences and so on. While I assume that such processes operate according to determinate physical principles, the system architecture that embodies them enables the ability to exercise the choice, flexibility and control that we experience — a form of biological determinism that is compatible with experienced free will.

So *who’s* in control? Who chooses, has thoughts, generates images and so on? We habitually think of ourselves as being our *conscious selves*. But it should be clear from the above that the different facets of our experienced, conscious selves are generated by and represent aspects of our own preconscious minds. That is, we are *both* the pre-conscious generating processes *and* the conscious results. Viewed from a third-person perspective our own preconscious mental processes look like neurochemical and associated physical activities in our brains. Viewed introspectively, from a first-person perspective, our preconscious mind seems like a personal but ‘empty space’ from which thoughts, images and feelings spontaneously arise. *We* are as much one thing as the other — and this requires a shift in our sensed ‘centre of gravity’ to one where our consciously experienced self becomes just the visible ‘tip’ of our own embedding, preconscious mind.

---

[19] Such functional differences are beyond the scope of this paper. However they have been extensively investigated, e.g. in studies of selective attention, controlled versus automatic processing, and so on (see e.g. Velmans, 1991; Kihlstrom, 1996).

## *APPENDIX*

### IS CONSCIOUSNESS NOTHING MORE THAN A STATE OF THE BRAIN?

It has long been suspected that there is a *causal relation* between mind or consciousness and brain. For example, Hippocrates of Cos (460–357 BC) wrote that,

Man ought to know that from the brain and from the brain only, arise our pleasures, joys, laughter and jests, as well as our sorrows, pains, griefs and fears. Through it, in particular, we think, see, hear, and distinguish the ugly from the beautiful, the bad from the good, the pleasant from the unpleasant, in some cases using custom as a test, in others perceiving them from their utility. It is the same thing which makes us mad or delirious, inspires us with dread and fear, whether by night or by day, brings sleeplessness, inopportune mistakes, aimless anxieties, absent-mindedness, and acts that are contrary to habit (from Jones, 1923; cited in Flew, 1978, p. 32).

However, the claim that mind or consciousness is *nothing more than* a state of the brain is far more radical. If this claim can be justified, then the fundamental puzzles surrounding the mind–body relationship, and (in its modern form) the consciousness–brain relationship would be solved. Clearly, if consciousness is nothing more than a state of the brain (a C-state say), it should be possible to understand it within the existing framework of natural science. Causal relations between consciousness and brain would translate into the causal relations between C-states and other brain states — and the functions of consciousness would simply be the functions of C-states within the global economy of the brain. The methods for investigating consciousness would then be third-person methods of the kind already well developed in neurophysiology and cognitive science. With such a potential prize in view, philosophical and scientific theories of consciousness over the last thirty years have in the main assumed, or tried to show, that some form of materialist reductionism is true.

#### **How could Conscious Experiences be Brain States?**

Given the apparent differences between the ‘qualia’ of conscious experiences and brain states it is by no means *obvious* that they are one and the same! Physicalists such as Ullin Place (1956) and J.J.C. Smart (1962) accepted that these apparent differences exist. They also accepted that descriptions of mental states and descriptions of their corresponding brain states are not identical in meaning. However, they claimed that with the advance of neurophysiology these descriptions will *be discovered* to be statements about one and the same thing. That is, a contingent rather than a logical identity will be established between consciousness, mind and brain.

Smart (1962) summarizes this position in the following way:

Let us first try to state more accurately the thesis that sensations are brain-processes. It is not the thesis that, for example, ‘after-image’ or ‘ache’ means the same as ‘brain-process of sort X’ (where ‘X’ is replaced by a description of a certain brain process). It is that, in so far as ‘after-image’ or ‘ache’ is a report of a process, it is a report of a process that happens to be a brain process. It follows that the thesis does

not claim that sensation statements can be translated into statements about brain processes. Nor does it claim that the logic of a sensation statement is the same as that of a brain process statement. All it claims is that in so far as a sensation statement is a report of something, that something is a brain process. *Sensations are nothing over and above brain processes* (p. 163 — my italics).

In short there is a distinction to be drawn between how things seem, how we describe them, and how they really are.

It is important to remember that no discovery that reduces consciousness to brain has yet been made. Physicalism, therefore, is partly an expression of faith, based on precedents in other areas of science — and arguments in defence of this position have focused on the *kinds of discovery that would need to be made* for reductionism to be true.

C.D. Broad noted in 1925 that materialism comes in three basic versions: *radical*, *reductive* and *emergent*. Radical materialism claims that the term ‘consciousness’ does not refer to anything real (in contemporary philosophy this position is usually called ‘eliminativism’). Reductive materialism accepts that consciousness does refer to something real, but science will discover that real thing to be nothing more than a state (or function) of the brain. Emergentism also accepts the reality of consciousness but claims it to be a higher-order property of brains; it supervenes on neural activity, but cannot be reduced to it.

While it is not the purpose of this Appendix to give a full appraisal of these positions (I do this elsewhere — Velmans, 2000, chs. 3, 4 and 5) it may be useful to indicate why I do not adopt them. So, by way of illustration, I list some of problems that physicalism must solve, some of the more plausible physicalist solutions to these, and a few of the problems with the solutions below.

### **What Non-Eliminative Reductionism Needs to Show**

Let us assume that, in some sense, our conscious experiences are real. To each and every one of us, our conscious experiences are observable *phenomena* (psychological *data*) which we can describe with varying degrees of accuracy in ordinary language. *Other* people’s experiences might be hypothetical constructs, as we cannot observe their experiences in the direct way that we can observe our own, but that does not make our own experiences similarly hypothetical. Nor are our own conscious experiences ‘theories’ or ‘folk psychologies’. We may have everyday theories *about* what we experience, and with deeper insight we might be able to improve them, but this would not replace or necessarily improve the experiences themselves.

In essence then, the claim that conscious experiences are nothing more than brain states is a claim about one set of phenomena (first-person experiences of love, hate, the smell of mown grass, the colour of a sunset, etc.) being nothing more than another set of phenomena (brain states, viewed from the perspective of an external observer). Given the extensive, apparent differences between conscious experiences and brain states this is a tall order. Formally, one must establish that despite appearances, conscious experiences are *ontologically identical* to brain states.



Instances where phenomena viewed from one perspective turned out to be one and the same as seemingly different phenomena viewed from another perspective do occur in the history of science. A classical example is the way the ‘morning star’ and the ‘evening star’ turned out to be identical (they were both found to be the planet Venus). But viewing consciousness from a first- versus a third-person perspective is very different to seeing the same planet in the morning or the evening. From a third-person (external observer’s) perspective one has *no direct access* to a subject’s conscious experience. Consequently, one has no third-person data (about the experience itself) which can be compared to or contrasted with the subject’s first-person data. Neurophysiological investigations are limited, in principle, to isolating the neural correlates or antecedent causes of given experiences. This would be a major scientific advance. But what would it tell us about the nature of consciousness itself?

### Common Reductionist Arguments and Fallacies

Reductionists commonly argue that if one can find the neural *causes* or *correlates* of consciousness in the brain, then this would establish consciousness *itself* to be a brain state (see for example, Place, 1956; Crick, 1994). Let us call these the ‘causation argument’ and the ‘correlation argument’. I suggest that such arguments are based on a fairly obvious fallacy. For consciousness to be nothing more than a brain state, it must be *ontologically identical* to a brain state. However, *correlation* and *causation* do not establish *ontological identity*. These relationships have been persistently confounded in the literature. So let me make the differences clear.

Ontological identity is *symmetrical*; that is, if A is identical to B, then B is identical to A. Ontological identity also *obeys Leibniz’s Law*: if A is identical to B, all the properties of A are also properties of B, and vice-versa (for example all the properties of the ‘morning star’ are also properties of the ‘evening star’).

Correlation is also *symmetrical*; if A correlates with B, then B correlates with A. But correlation does *not obey Leibniz’s Law*; if A correlates with B, it does not follow that all the properties of A and B are the same. For example, height in humans correlates with weight, but height and weight do not have the same set of properties.

Causation, by contrast, is *asymmetrical*; if A causes B, it does not follow that B causes A. If a rock thrown in a pond causes ripples in the water, it does not follow that ripples in the water cause the rock to be thrown in the pond. And causation does *not obey Leibniz’s Law* (flying rocks and pond ripples have very different properties).

Once the obvious differences between causation, correlation and ontological identity are laid bare the weaknesses of the ‘causation argument’ and the ‘correlation argument’ are clear. Under appropriate conditions, brain states may be shown to cause, or correlate with conscious experiences, but it does not follow that conscious experiences are nothing more than states (or, for that matter, functions) of the brain. To demonstrate that, one would have to establish an ontological identity in which all the properties of a conscious experience and

corresponding brain state are identical. Unfortunately for reductionism, few if any properties of experiences (accurately described) and brain states appear to be identical.

In short, the causes and correlates of conscious experience should not be confused with their *ontology*. As it happens, various *nonreductionist* positions such as dualist-interactionism, epiphenomenalism and modern dual-aspect theory *agree* that consciousness (in humans) is causally influenced by and correlates with neural events, but they *deny* that consciousness is nothing more than a state of the brain. As no information about consciousness *other than its neural causes and correlates* is available to neurophysiological investigation of the brain, it is difficult to see how such research could ever settle the issue. The *only* evidence about what conscious experiences are like comes from first-person sources, which consistently suggest consciousness to be something other than or additional to neuronal activity. Given this, I conclude that reductionism via this route *cannot be made to work* (cf. Velmans, 1998).

### **False Analogies**

Faced with this difficulty, reductionists usually turn to analogies from other areas in science, where a reductive, causal account of a phenomenon led to an understanding of its ontology, very different to its phenomenology. Francis Crick (1994), for example, makes the point that in science, reductionism is both common and successful. Genes for example turned out to be nothing but DNA molecules. So, in science, this is the best way to proceed. While he recognizes that experienced (first-person) ‘qualia’ pose a problem for reductionism, he suggests that in the fullness of time it may be possible to describe the *neural correlates* of such qualia. If we can understand the nature of the correlates we may come to understand the corresponding forms of consciousness. By these means science will show that ‘You’re nothing but a pack of neurones!’.

It should be apparent from the above that finding the neural correlates of consciousness won’t be enough to reduce people to neurones! The reduction of consciousness to brain is also quite unlike the reduction of genes to DNA. In the development of genetics, ‘genes’ were initially hypothetical entities inferred to exist to account for observed regularities in the transmission of characteristics from parents to offspring. The discovery that genes are DNA molecules shows how a theoretical entity is sometimes discovered to be ‘real’. A similar discovery was made for bacteria, which were inferred causes of disease until the development of the microscope, after which they could be seen. Viruses remained hypothetical until the development of the electron microscope, after which they too could be seen. These are genuine cases of materialist reduction (of hypothetical to physical entities).

But it would be absurd to regard conscious experiences as ‘hypothetical entities’, waiting for their neural substrates to be discovered to make them real. Conscious experiences are first-person *phenomena*. To those who have them, they provide the very fabric of subjective reality. One does not have to wait for the advance of neuroscience to know that one has been stung by a bee! If conscious

experiences *were* merely hypothetical, the mind–body problems, and in particular the problems posed by the phenomenal properties of ‘qualia’ would not exist.

Ullin Place (1956) focuses on causation rather than correlation. As he notes, we now understand lightning to be nothing more than the motion of electrical charges through the atmosphere. But mere correlations of lightning with electrical discharges do not suffice to justify this reduction. Rather, he argues, the reduction is justified once we know that the motion of electrical charges through the atmosphere *causes* what we experience as lightning. Similarly, a conscious experience may be said to be a given state of the brain once we know that brain state to have *caused* the conscious experience.

I have dealt with the fallacy of the ‘causation argument’ above. But the lightning analogy is seductive because it is half-true. That is, *for the purposes of physics* it is true that lightning can be described as nothing more than the motion of electrical charges. But there are three things that need to be accounted for in this situation, not just one — an event in the world, a perceiver, and a resulting experience. Physics is interested in the nature of the event in the world. However, psychology is interested in how this physical event interacts with a visual system to produce *experienced lightning* — in the form of a perceived flash of light in a phenomenal world. This experienced lightning may be said to *represent* the same event in the world which physics describes as a motion of electrical charges. But the *phenomenology of the experience itself* cannot be said to be nothing more than the motion of electrical charges! Prior to the emergence of life forms with visual systems on this planet, there presumably was no such phenomenology, although the electrical charges which now give rise to this experience did exist.

In sum, the fact that motions of electrical charges cause the experience of lightning does not warrant the conclusion that the *phenomenology* of the experience is nothing more than the motion of electrical charges. Nor would finding the neurophysiological causes of conscious experiences warrant the reduction of the phenomenology of those experiences to states of the brain.

Given that examples of first-person reduction to third-person science (DNA, lightning, colour, heat, etc.) are not really examples of first-person reduction at all, perhaps a nonreductive materialism is more appropriate. For example, according to Searle (1987; 1992; 1994a; 1997) conscious states cannot be redescribed (now or ever) in neurophysiological language. Rather, they have to be described just as they seem to be. Searle, for example, believes *subjectivity* and *intentionality* to be essential features of consciousness. Conscious states have ‘intrinsic intentionality’, that is, it is intrinsic to them that they are *about* something. According to Searle, this distinguishes conscious states from physical representations such as sentences written on a page. Conscious readers might interpret these *as if* they are about something (such physical representations have ‘as-if intentionality’), but they are just marks on a piece of paper and not about anything in themselves. Subjectivity, too, ‘is unlike anything else in biology, and in a sense it is one of the most amazing features of nature’ (Searle 1994a, p. 97). Nevertheless, he maintains that conscious states are just higher-order features of the brain.

### Emergentism

In classical dualism, consciousness is thought to be a nonmaterial substance or entity different in kind from the material world, with an existence that is independent of the existence of the brain (although in normal life it interacts with the brain). ‘Emergentism’ in the form of ‘property dualism’ retains the view that there are fundamental differences between consciousness and physical matter, but views these as different kinds of property of the brain. That is, consciousness is not *reducible* but its existence is still *dependent* on the workings of the brain — and according to Searle, such a non-reducible brain property is still ‘physical’.

Searle (1987), for example, argues that *causality* should not be confused with *ontological identity* (as I do in my critique of reductionism above), and his case for physicalism appears to be one of the few to have addressed this distinction head-on. The gap between what *causes* consciousness and what conscious *is* can be bridged, he suggests, by an understanding of how microproperties relate to macroproperties. Liquidity of water is caused by the way H<sub>2</sub>O molecules slide over each other, but is nothing more than (an emergent property of) the combined effect of these molecular movements. Likewise, solidity is caused by the way molecules in crystal lattices bind to each other, but is nothing more than the higher order (emergent) effect of such bindings. In similar fashion, consciousness is caused by neuronal activity in the brain and is nothing more than the higher order, emergent effect of such activity. That is, consciousness is just a *physical macroproperty* of the brain.

Searle’s argument is attractive, but it needs to be examined with care. The brain undoubtedly has physical macroproperties of many kinds. Like other physical systems, its physical microstructure supports a physical macrostructure. However, the physical macroproperty of brains that is most closely analogous to ‘solidity’ and ‘liquidity’ is ‘sponginess’, not consciousness! There are, of course, more psychologically relevant macroproperties, for example the blood flow patterns picked up by PET scans or the magnetic and electrical activities detected by fMRI and EEG. But why should increased blood flow constitute subjectivity, or why would it be ‘like anything’ to be an electrical potential or magnetic field? While some of these properties undoubtedly *correlate* with conscious experiences, there is little reason to suppose that they are *ontologically identical* to conscious experiences.

One might also question how Searle’s property dualism could really be a form of *physicalism*. Searle insists that consciousness is a *physical* phenomenon, produced by the brain in the sense that the gall bladder produces bile. But he also stresses that *subjectivity* and *intentionality* are defining characteristics of consciousness. Unlike physical phenomena, the phenomenology of consciousness cannot be observed from the outside; unlike physical phenomena, it is always *of* or *about* something. So, even one accepts that consciousness is, in some sense, caused by or emergent from the brain, why call it ‘physical’ as opposed to ‘mental’ or ‘psychological’? Merely *relabelling* consciousness, or moving from

micro- to macroproperties doesn't really close the gap between 'objective' brains and 'subjective' experiences!<sup>20</sup>

In sum, demonstrating the brain to have physical macroproperties that are supervenient on its physical microproperties is one thing; *identifying* those physical macroproperties with the properties of *consciousness* is another! Searle, as shown above, tries to settle the issue by *fiat*. Subjective, intentional conscious experiences are simply *declared* to be physical states. But this doesn't really help much. The ontology of these 'new' physical states is not really clarified by renaming them. Nor does the transition from smaller things to larger things (from microproperties to macroproperties) really *explain* how material brains viewed from a third-person perspective could themselves have a conscious, first-person perspective! Further, the problem of *how* such extraordinary 'subjective', 'intentional' states could *interact* with ordinary physical states remains. (A fuller analysis of Searle's position, taking account of his 1997 defence, is given in Velmans, 2000, ch. 3.)

---

[20] I should stress that I do not deny that conscious experiences can be said to 'emerge' from the human brain in the sense that given brain states can be said to *cause* given conscious experiences. That is, I do not deny the legitimacy of physical→mental causal accounts, anymore than I deny the legitimacy of physical→physical, mental→physical and mental→mental accounts. The question is: how do we make sense of these accounts? The physicalist answer (in whatever guise it takes) is to translate all these causal accounts into physical→physical accounts — in this case, by trying to show that conscious states are nothing more than higher-order, emergent *physical* states of the brain. As far as I can tell, this manoeuvre cannot really be made to work. That is, first-person consciousness cannot be thought of as a 'physical' property of the brain in any conventional, third-person sense of the term 'physical'. Note that the problems of *identifying* first-person consciousness with third-person features persist even when we select plausible, emergent brain properties that are less obviously 'physical', but nevertheless describable in third-person, functional terms. For example, Dewar (1976) (elaborating on the emergent-interactionism of Roger Sperry, 1969) cites the phenomenon of 'mutual entrainment'. The term 'entrainment' refers to the synchronization of an oscillator to an input signal. This occurs, for example, when television receiver oscillators controlling the vertical and horizontal lines 'lock into' transmitting frequencies to produce a given picture on the screen. Examples of entrainment, Dewar notes, may also be found at many levels of biological organization — a particularly apposite case being the way 'biological clocks' governing circadian rhythms can be locked into varying periods (of around 24 hours) to produce altered cycles of day–night activity in animals. 'Mutual entrainment' occurs when two or more oscillators interact in such a way that they pull one another into synchrony. This occurs, for example, when different alternating-current generators feeding the national grid are pulled into synchrony by what Norbert Wiener refers to as a 'virtual governor' in the system. Although the generators may be far distant from each other and may start up and stop at idiosyncratic times, once 'on-line' they are made to speed up or slow down to produce AC current in phase with that of all the other machines feeding the grid. As Dewar points out the 'virtual governor' is not located in any one place in the system, but rather pervades the system as a whole so that it does not have a 'physical existence' in the usual sense. It is an emergent property of the entire system. In similar fashion, Dewar suggests, consciousness is 'a holistic emergent property of the interaction of neurones which has the power to be self-reflective and ascertain its own awareness'.

This analogy becomes particularly interesting in the light of the suggestion that synchronous or correlated firing of diverse neurone groups (at rhythmic frequencies in the 40 Hz region) might produce the 'neural binding' required to produce an integrated experience from features of objects that are encoded in spatially separated regions of the brain. Given the well-integrated nature of normal conscious experiences, it seems reasonable to propose that binding processes operate prior to the formation of, or co-occur with, such experiences. However, there is little reason to suggest that 'binding' or 'mutual entrainment' is *ontologically identical* to consciousness — unless we are willing to accept that the national grid is conscious. How mutual entrainment or binding 'has the power to be self-reflective and ascertain its own awareness' remains a mystery! (A more detailed analysis of how consciousness relates to mutual entrainment and binding is given in Velmans, 2000, pp. 41–2).



## References

- Baars, B.J. and McGovern, K. (1996), 'Cognitive views of consciousness: What are the facts? How can we explain them?', in Velmans (1996a).
- Barber, T.X. (1984), 'Changing "unchangeable" bodily processes by (hypnotic) suggestions: a new look at hypnosis, cognitions, imagining, and the mind body problem', in *Imagination and Healing*, ed. A.A. Sheikh (Farmingdale, NY: Bayworld).
- Boomer, D.S. (1970), 'Review of F. Goldman-Eisler *Psycholinguistics: Experiments in Spontaneous Speech*', *Lingua*, **25**, pp. 152–64.
- Broad, C.D. (1925), *The Mind and Its Place in Nature* (London: Routledge & Kegan Paul).
- Chalmers, D.J. (1996), *The Conscious Mind* (New York: Oxford University Press).
- Crick, F. (1994), *The Astonishing Hypothesis* (London: Simon & Schuster).
- Dewar, E.M. (1976), 'Consciousness in control systems theory', in *Consciousness and the Brain*, ed. G.G. Globus, G. Maxwell and I. Savodnik (New York: Plenum).
- Flew, A. (ed. 1978), *Body, Mind, and Death* (New York: Macmillan Publishing).
- Fodor, J.A., Bever, T.G. and Garrett, M.F. (1974), *The Psychology of Language* (New York: McGraw-Hill).
- Gardner, H. (1987), *The Mind's New Science* (New York: Basic Books).
- Goldman-Eisler, F. (1968), *Psycholinguistics* (New York: Academic Press).
- Hashish, I., Finman, C. and Harvey, W. (1988), 'Reduction of postoperative pain and swelling by ultrasound: a placebo effect', *Pain*, **83**, pp. 303–11.
- Kanttinen, N. and Lyytinen, H. (1993), 'Brain slow waves preceding time-locked visuo-motor performance', *Journal of Sport Sciences*, **11**, pp. 257–66.
- Karrer, R., Warren, C. and Ruth, R. (1978), 'Slow potentials of the brain preceding cued and non-cued movement: effects of development and retardation', in *Multidisciplinary Perspectives in Event-Related Potential Research*, ed. D.A. Otto (Washington DC: US Government Printing Office).
- Kihlstrom, J.F. (1996), 'Perception without awareness of what is perceived, learning without awareness of what is learned', in Velmans (1996a).
- Kornhuber, H.H. and Deecke, L. (1965), 'Hirnpotentialänderungen bei willkürbewegungen und passiven bewegungen des menschen: Bereitschaftspotential und reafferente potentiale', *Pflügers Archiv für die Gesamte Physiologie des Menschen und Tiere*, **284**, pp. 1–17.
- Lenneberg, E.H. (1967), *Biological Foundations of Language* (New York: Wiley).
- Libet, B. (1985), 'Unconscious cerebral initiative and the role of conscious will in voluntary action', *Behavioral and Brain Sciences*, **8**, pp. 529–66.
- Libet, B. (1996), 'Neural processes in the production of conscious experience', in Velmans (1996a).
- Libet, B., Wright Jr., E.W., Feinstein, B. and Pearl, D.K. (1979), 'Subjective referral of the timing for a conscious experience: A functional role for the somatosensory specific projection system in man', *Brain*, **102**, pp. 193–224.
- Mangan, B. (1993), 'Taking phenomenology seriously: The 'fringe' and its implications for cognitive research', *Consciousness and Cognition*, **2** (2), pp. 89–108.
- McMahon, C.E. and Sheikh, A. (1989), 'Psychosomatic illness: a new look', in *Eastern and Western Approaches to Healing*, ed. A. Sheikh and K. Sheikh (New York: Wiley-Interscience).
- Pelletier, K.R. (1993), 'Between mind and body: stress, emotions, and health', in *Mind Body Medicine*, ed. D. Goleman and J. Gurin (New York: Consumer Reports Books).
- Place, U. (1956), 'Is consciousness a brain process?', *British Journal of Psychology*, **47**, pp. 44–50.
- Rakover, S.S. (1996), 'The place of consciousness in the information processing approach: The mental pool thought experiment', *Behavioral and Brain Sciences*, **19** (3), pp. 535–6.
- Searle, J. (1987), 'Minds and brains without programs', in *Mindwaves*, ed. C. Blakemore and S. Greenfield (Oxford: Blackwell).
- Searle, J. (1992), *The Rediscovery of the Mind* (Cambridge, MA: MIT Press).
- Searle, J. (1994a), 'The problem of consciousness', in *Consciousness in Philosophy and Cognitive Neuroscience*, ed. A. Revonsuo and M. Kamppinen (Hillsdale, NJ: Lawrence Erlbaum Assoc.).
- Searle, J. (1994b), 'Intentionality (1)', in *A Companion to the Philosophy of Mind*, ed. S. Guttenplan (Oxford: Blackwell).
- Searle, J. (1997), *The Mystery of Consciousness* (London: Granta Books).



- Smart, J.J.C. (1962), 'Sensations and brain processes', in *Philosophy of Mind*, ed. V.C. Chappell (Englewood Cliffs: Prentice-Hall).
- Sheikh, A.A. (ed. 2001), *Healing Images: The Role of Imagination in the Healing Process* (Amityville, NY: Baywood Publishing Company).
- Sheikh, A.A., Kundendorf, R.G. and Sheikh, K.S. (1996), 'Somatic consequences of consciousness', in Velmans (1996a).
- Skrabaneck, P. & McCormick, J. (1989), *Follies and Fallacies in Medicine* (Glasgow: The Tarragon Press).
- Sperry, R.W. (1969), 'A modified concept of consciousness', *Psychological Review*, **76** (6), pp. 532–6.
- Stoffregen, T.A. and Benoit, G.B. (2001), 'On specification of the senses', *Behavioral and Brain Sciences*, **24** (2), pp. 195–261.
- Syrjala, K.A. and Abrams, J.R. (1996), 'Hypnosis and imagery in the treatment of pain', in *Psychological Approaches to Pain Management: A Practitioner's Handbook*, ed. R.J. Catchel and D.C. Turk (New York: The Guildford Press).
- Tye, M. (1995), *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind* (Cambridge, MA: MIT Press).
- Velmans, M. (1990), 'Consciousness, brain, and the physical world', *Philosophical Psychology*, **3**, pp. 77–99.
- Velmans, M. (1991a), 'Is human information processing conscious?', *Behavioral and Brain Sciences*, **14** (4), pp. 651–69.
- Velmans, M. (1991b), 'Consciousness from a first-person perspective', *Behavioral and Brain Sciences*, **14** (4), pp. 702–26.
- Velmans, M. (1993), 'Consciousness, causality and complementarity', *Behavioral and Brain Sciences*, **16** (2), pp. 409–16.
- Velmans, M. (ed. 1996a), *The Science of Consciousness: Psychological, Neuropsychological and Clinical Reviews* (London: Routledge).
- Velmans, M. (1996b), 'Consciousness and the "causal paradox"', *Behavioral and Brain Sciences*, **19** (3), pp. 538–42.
- Velmans, M. (1998), 'Goodbye to reductionism', in *Towards a Science of Consciousness II: The Second Tucson Discussions and Debates*, ed. S. Hameroff, A. Kaszniak and A. Scott (Cambridge, MA: MIT Press).
- Velmans, M. (2000), *Understanding Consciousness* (London: Routledge/Psychology Press).
- Velmans, M. (2001a), 'A natural account of phenomenal consciousness', *Consciousness and Communication*, **34** (1&2), pp. 39–59.
- Velmans, M. (2001b), 'Heterophenomenology versus critical phenomenology: A dialogue with Dan Dennett', <http://cogprints.soton.ac.uk/documents/disk0/00/00/17/95/index.html>
- Wall, P.D. (1996), 'The placebo effect', in *The Science of Consciousness: Psychological, Neuropsychological and Clinical Reviews*, ed. M. Velmans (London: Routledge).
- Watkins, A. (1997), 'Mind–body pathways', in *Mind–Body Medicine: A Clinician's Guide to Psychoneuroimmunology*, ed. A. Watkins (New York: Churchill Livingstone).
- Watt, D. (2000), 'The centrencephalon and thalamocortical integration: Neglected contributions of periaqueductal gray', *Consciousness and Emotion*, **1** (1), pp. 91–11.