The Problem of Consciousness and the Innerness of the Mind

Jim Hopkins

1. Consciousness and innerness

The problem of consciousness is commonly taken to be that of accounting for aspects of experience which we describe as *inner*, and also as phenomenal, subjective and private. Thus for example David Chalmers introduces what he calls 'the real mystery of the mind' as follows:

> From an *objective* viewpoint, the brain is relatively comprehensible. When you look at this page, there is a whirr of processing: photons strike your retina, electrical signals are passed up your optic nerve and between different areas of your brain … . But there is also a *subjective* aspect. When you look at the page you are conscious of it, directly experiencing the images and words as part of your *private*, mental life .... The hard problem...is the question of how physical processes in the brain give rise to *subjective* experience. This puzzle involves *the inner aspect* of thought and perception ... . It is these phenomena which constitute the real mystery of the mind. (Chalmers 1995, p. 92)

In speaking of 'the inner aspect of thought and perception' Chalmers draws on a familiar feature of our everyday conception of the mind. We speak of the mind as a locus which things are *in*, and these things as therefore having an *internal* aspect; likewise we speak of ourselves as aware of things *from the inside*, or by *introspection*; of our *inner lives*; and so on.

We seem to speak and think in this way from early childhood, and to distinguish things which we take to be 'in the mind' from the physical things we can see or touch. Thus as Henry Wellman writes

> With regard to ontology, children as young as three years firmly divide the mental and physical worlds … if told about someone who has a dog that has run away, versus someone who is thinking of a dog, three-year-olds know that neither 'dog' can be seen or petted, but that none the less one is mental ('just in his mind', 'only imagination') whereas the other is physically real, and merely unavailable. In these ways young children appropriately distinguish between real and mental entities. (Wellman 1993, p. 14)

Our notion of things as internal to the mind thus seems to go with a natural intuitive dualism. Philosophical writers have developed this idea in many ways, often speaking as if the mind were a kind of inner *place* or *space*, in which mental phenomena are somehow housed. Augustine, for example, stressed how as a baby he found 'my wants were *within me* and those others [who might satisfy them] were outside; nor had they any faculty enabling them to *enter* my mind. So I would fling my arms and legs about and utter signs ... .' (Augustine, p. 6) And we can readily skip the centuries to comparable recent discussions, such as the following by Thomas Metzinger.

> To be able to speak seriously about a *science of consciousness*, a number of fundamental questions would have to be answered. It is interesting to note that with the emergence of consciousness private worlds – spaces of inner experiences – are opened up. These spaces, however, are *individual* spaces: ego-centres of experience that suddenly appear in a centerless universe. Each such centre of consciousness constitutes its own perspective on the world. This perspective is what philosophers sometimes like to call the 'first-person perspective'. A phenomenal world of its own is tied to each of these perspectives. These individual worlds of experience also possess a historical dimension: almost always a psychological biography emerges together with them – what we call our 'inner life'. This too can be seen as the history of the genesis of a world, or a *phenomenal* cosmology: within each of us a cosmos of consciousness unfolds temporarily, a *subjective* universe develops. The first part of the problem is to understand how a variety of subjective universes can constantly form and disappear in our objective universe ... . (Metzinger 1995, p. 5)

Widespread and apparently familiar as they are, these ideas of innerness ought surely to be regarded as puzzling. What does it mean to say that something exists *within* the mind, if this is to be contrasted with existence which is physical? And how does this relate to the idea of things, including the mind, being *inside* the physical body? In what follows I propose an approach to these questions, and try to show how it casts light on the problem of consciousness.

2. The double internality of the mental

It is noteworthy that on this way of thinking mental phenomena are ascribed a *double innerness.* They are said to go on *within* the mind, and this in turn to be located *within* the body. To bring this out more fully we can illustrate it with a kind of cartoon, as follows:

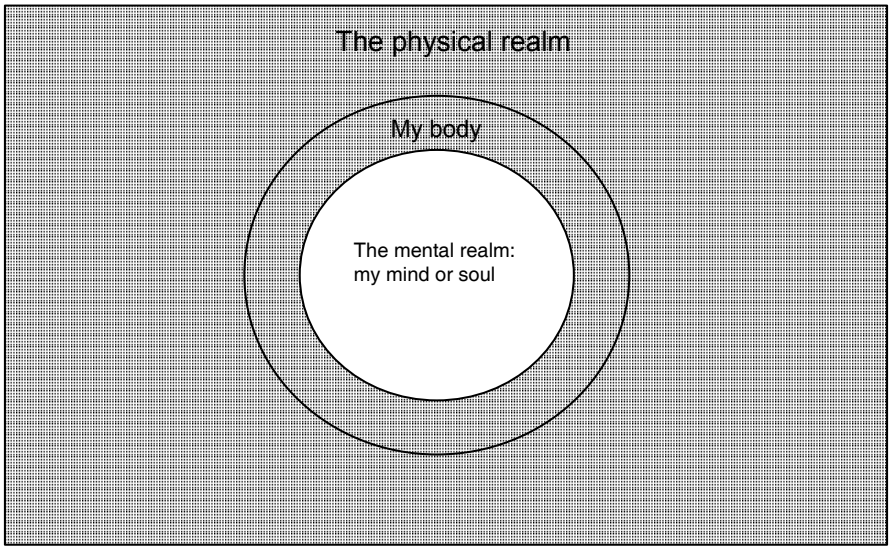

Now when we put the matter this graphically, we may wonder how far, despite speaking in this way, we actually do entertain such an image of ourselves. This is a serious question, and we often reject what Wittgenstein called the 'pictures' underlying our uses of words when we make them more explicit.[1] We can gain a more vivid sense of this by considering the phenomena associated with a phantom limb. The bearer of a phantom arm, like a person with a real arm, feels himself to have an arm which is a locus of sensation, willing, etc. In the case of the phantom arm, however, the felt locus is actually empty, and can be seen to be so. The feelings connected with the phantom thus seem *distributed in the physical space where the arm would be*, but also *not to be real occupants of this space*, in the sense of *things which actually have extension within it*. Rather, and in accord with the conception we are considering, the feelings seem located in a *subjective inner space* – that of the bearer's mind – which is encompassed by, but somehow distinct from, the physical space which surrounds it.

Vilayanur Ramachandran has devised a way of enabling patients with phantom arms to have an experience as of seeing the phantom as well as feeling it. This is done using a box with holes in the sides and an internal mirror. The patient puts his intact arm in one of the holes and puts the phantom (or imagines putting it) into the other. The mirror is so placed as to allow the patient to see his intact arm, but also to reflect the image of that

[1] This is the 'picture' of the mind which Wittgenstein discusses often in *Philosophical Investigations*, and caricatures as that of the 'beetle in the box'. This is discussed further in my 2000.

arm, so that it also appears (and in mirror-image right-to-left reversal) as in the place where the phantom has been 'put', and feels to be. When the patient looks into the box, he sees himself as having two intact arms, in just the places where he feels these arms to be. In effect his normal image of his missing arm – as having feelings going on within it, which are also encompassed by his physical body – is temporarily (and illusorily) regained. The phantom arm which he previously felt *as from inside* but could not see now appears *as visible*, and hence *from outside*. If asked to move both arms or hands in symmetrical ways, the patient may also have the experience both of moving the phantom and seeing the arm itself move accordingly.

This, as Ramachandran notes, may enable the patient to manipulate the phantom as if it were real. He describes a patient who suffered discomfort from the apparent position of a phantom arm, which he had not been able to alter.

> Philip rotated his body, shifting his shoulder to 'insert' his lifeless phantom in the box. Then he put his right hand on the other side of the mirror and attempted to make synchronous movements. As he gazed into the mirror he gasped and then cried out, 'Oh my God! Oh my God doctor! This is unbelievable. Its mind-boggling!' He was jumping up and down like a kid. 'My left arm is plugged in again… I can move my arm again. I can feel my elbow moving, my wrist moving. Its all moving again.' (Ramachandran 1998, p. 8)

In considering first the phantom limb and then this illusion of seeing it, we as it were leave and return to normal experience by rotating through two right angles. Doing this may remind us how apt it is to regard mental phenomena as both distinct from our bodies and yet also housed inside them. For this seems particularly clearly to be the nature of normal experience as it is recovered by the bearer of the phantom limb in Ramachandran's box. His experience of himself *from inside* includes that of his arm, as a place of feeling and potential exercise of will. But although such a place is clearly there in his mind, it is 'just in his mind'. For as the loss of his physical arm makes clear, this inner locus is not really part of the physical space in which extended things are to be found. Once placed in Ramachandran's box, however, this psychological but non-physical place is re-clothed in flesh, and so again seems fully real. The space of the arm, as felt from inside, now seems integrated with a bodily arm which is subject to its owners will. The situation seems restored to that which we describe in our everyday language of the inner and the outer, and which is schematically depicted in the drawing above. So something like this image

seems really to represent an aspect of our experience of mental phenomena, even if we rarely make it fully explicit, and are surprised enough by its implications to be tempted to withdraw it when we do so.

3. The inner aspects of the mental as phenomenal, subjective, and private

This example also indicates the ways we contrast the phenomenal characteristics of mental events with those of external physical things. Something *physical* and *external* to the mind, such as a tree, or again the flesh of an arm, can be *externally perceived*. Such things can be seen, touched, or otherwise made the object of the senses, and hence located in the encompassing physical space in which we find ourselves. Accordingly they have physical properties such as size, shape, and material composition, and admit further description in terms of the physical sciences, such as cellular and chemical structure, mass, weight, energy, and so on.

So also such things are *objective*. They exist independently of our perceiving or thinking about them, and they have a real nature distinct from what is apparent in our perception or understanding. Again, and connectedly, such things are *public*, in the sense that they are *capable of being perceived and investigated by more than one person*. These features, moreover, seem intrinsic and interconnected. The world which we think of as *external* to the mind and *perceivable by the senses* is the *physical* world, and it is in the nature of what is physical to be *objective* and *public* in these or comparable ways.

This, as it seems, stands in sharp contrast to the way think of experiences such as the feeling of pain, or the visual experience of seeing a tree. Since these things exist *within* the mind, they cannot be seen by the eyes, touched, smelled, etc.; rather they are presented in *introspection*. Their being in the mind, moreover, is not (or not just) a matter of their being located within the physical space of the body, or in the brain or nervous system. Rather their innerness seems to be of an entirely different kind. We may be inclined to locate visual sensations, say, as just behind the eyes, and we may think of them as disposed in a spatial way, for example as constituting a *visual field.* But we do not think that if we physically examine the space behind the eyes we will see, or otherwise detect, this spatial field. As with the feelings of the phantom arm, visual experience and the visual field seem not to be in the same space as the retina, the optic nerve, etc.

Something similar applies in the case of pains, which have a precise bodily location. We feel the pain of an aching tooth as *in the tooth*, but we

also hold that no external examination of the space occupied by the tooth – no x-ray, scanning process, or the like – will reveal the pain itself which we feel there. If I touch an aching tooth with the tip of my tongue – even if I manage to put my tongue exactly at the place where the ache is felt to be – the tip of the tongue and the pain still seem to be in different places, which as it were nearly overlap. The felt quality of the pain, unlike the tip of the tongue, seems to be in a purely internal locus, comparable to that we noted in connection with the phantom limb, and accessible only to the person who actually has the toothache. Thus, as we say, these things are *in the mind*; but unlike the nervous system or the brain they apparently cannot be seen or otherwise detected *in the body* or *in the head*. Like the sensations and exercise of will relating to the phantom limb, they seem in, and part of, the *self*, but not *part of the physical body*, even though, like the sensations in the non-existent phantom limb, they seem distributed in space.

The inner aspects of sensations, experiences, and other conscious mental items are often called *phenomenal* qualities, or *qualia*. To many it seems that these phenomenal qualities are *non-physical*, in the sense that it is unintelligible or inexplicable that they should be the qualities of a physical thing. Again, these inner qualities seem *subjective*, in the sense (and by contrast to the physical things considered above) that they seem to *depend for their existence or nature upon the person who apprehends them*, and *to reveal their whole nature in this subjective apprehension*. When I feel a pain, it seems this pain could not exist if I were not feeling it – if I were to cease to feel the pain, it seems, the pain itself would cease to be, or cease to be painful. Likewise I seem to be acquainted with the whole inner nature of the pain. Unlike a physical object which I might see with my eyes, a pain has no inside which is concealed by its surface, nor any backside which is invisible to me. And since there is no more to the pain I am feeling than I now feel, I take the pain actually to be as I feel it to be.[2]

---

[2] This notion of subjectivity was well expressed by Thomas Reid:

> When I am pained, I cannot say that the pain I feel is one thing, and that my feeling it is another thing. They are one and the same thing, and cannot be disjoined, even in imagination. Pain, when it is not felt, has no existence. It can be neither greater nor less in degree nor duration, nor anything else in kind than it is felt to be. It cannot exist by itself, nor in any subject but a sentient being. No quality of an inanimate being can have the least resemblance to it. What we have said of pain may be applied to any other sensation. The feeling and the thing felt are one and the same. (Reid 1785, 1.12.)

Also, and connectedly, the inner quality of a particular pain seems *private*, in the sense that it can be felt or apprehended only by the person who has the pain. When I am aware of the inner nature of a pain, it seems that no one else can be aware of the inner nature of that particular pain; for, as noted above, my awareness seems to be of an internal locus which is available only to me (and the same holds for the sensations pertaining to a phantom limb). These features of experience again seem both intrinsic and interconnected; for it is the *inner* (introspectible) aspects of experience which we take to be *phenomenal*, and these phenomenal aspects which are *subjective* and *private*.

4. The inner and physicalism

These considerations provide a more fully developed account of the intuitive distinction, which we apparently make from early childhood, as between things which are inside as opposed to outside the mind. The difference is marked by what we can regard as a connected and perhaps overlapping series of oppositions, including those between

(i) Inner or internal to the mind as opposed to outer or external to it.
(ii) Introspectible (internally perceivable) as opposed to visible, tangible, etc. (externally perceivable)
(iii) Phenomenal as opposed to physical.
(iv) Subjective as opposed to objective.
(v) Private as opposed to public.

We can illustrate these oppositions by redrawing and extending our initial diagram as follows:

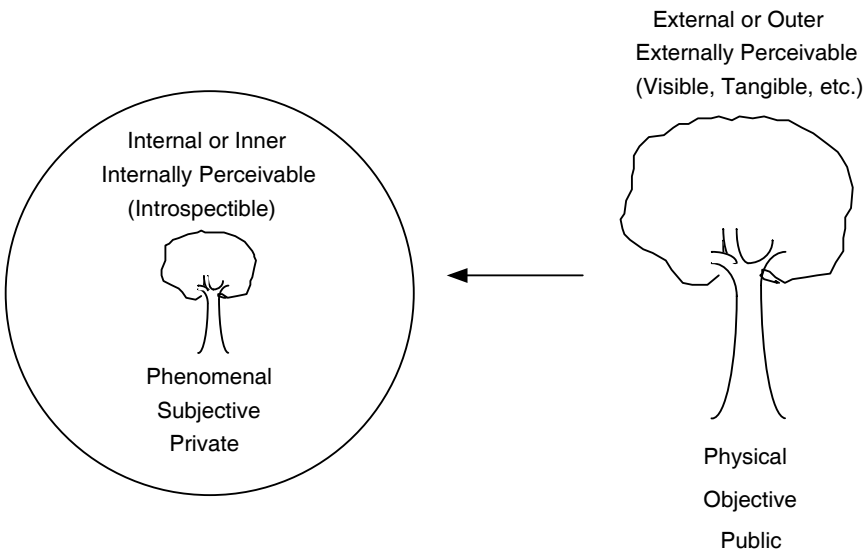

The internality of the mental

As this illustrates, the oppositions we are considering all seem rooted in that between the internal and the external. We conceive the mind as in the body, and mental phenomena, again, as occurring within the mind. Introspection is the internal analogue of perception, and phenomenal properties are those manifest in the already inner field of the mind. The subjectivity and privacy of these properties, therefore, seems consequent on their being internally perceivable by just one person, and from just one point of view. The objectivity and publicity of physical things, by contrast, seems secured by their existing in an external public space, and so being perceivable by more than one person and from more than one point of view.

These oppositions are party specified via negation, and they seem contradictory. It seems that *no single thing* could be both inner and outer, introspectible and externally visible (or tangible), phenomenal and physical, objective and subjective, or private and public. Hence, as Nagel, Chalmers, Metzinger and very many others have stressed, these internal aspects of mental phenomena seem particularly difficult to understand in physical terms, and give rise to the problem of consciousness.

5. The inner and the problem of consciousness

(a) Two aspects of the problem: reconciliation and derivation

Taking the problem of consciousness in this perspective, we can regard it as having two connected aspects. The first we may describe as a problem

of *reconciliation*, and the second as one of *derivation*. To approach the problem it is important to distinguish them clearly.

Reconciliation: To reconcile things, in the sense involved here, is to show how they are compatible. So as regards consciousness, the problem of reconciliation is that of showing how the qualities we assign to conscious events (particularly those of being inner, phenomenal, subjective, and private) are *compatible* with those we assign to neural events (particularly those of being outer, physical, objective, and public). Nagel, for example, expresses this problem by saying that 'We do not at present possess the conceptual equipment to understand *how subjective and physical features could both be essential aspects of a single entity or process*.' (Nagel 2000, p. 444)

As Nagel notes, this problem is not solved by maintaining that conscious events are in fact events in the brain, for the difficulty is to understand *how* this can be so – how the properties which we ascribe to conscious events are compatible with those which we ascribe to physical ones.[3] Hence it is also clearly not solved by holding that neural events, while physical, also have aspects or properties which are inner, phenomenal, subjective, or private, for this again is a version of what requires to be explained.

Derivation: The problem of derivation is related to that of reconciliation, but goes beyond it. This problem is raised by considering the kinds of explanation we attain in the physical sciences. In general physical science enables us to understand how the manifest features of things can be *derived from* their physical nature, and so to exhibit how they are consequent on it. So if physicalism is true, it seems that this should hold for the manifest inner features of experience as well.

If we describe the problem of opposition as that of seeing how certain physical events *could possibly have* the features of conscious events, or vice-versa, then we can describe the problem of derivation as that of seeing how the physical events in question *ought certainly to have* the mental features of conscious events. This is the problem of seeing how the inner, phenomenal, subjective, and private aspects of these events actually *flow from*, or *are consequence of*, their physical nature.

[3] This of course was stressed by Nagel in his celebrated Nagel (1979).

(b) Derivation and explanation.

We can spell this out more fully as follows. The natural sciences enable us to explain the manifest features of the physical world. We see the apparent movements of the stars across the night sky, and we have come to be able to explain these movements. We observe that ice and iron are solid, that water and glass are transparent, and we can explain why these things have these features. In each such explanation there is some form of *derivation* of the manifest and observable properties we explain, from other and more basic properties, which we find or hypothesize in the course of scientific investigation. In some cases we can cast these derivations in the form of rigorous logical or mathematical arguments which proceed from a few simple laws; and in others we are inclined to hold that such derivations are available in principle, even though they are too complex, or require too much information or calculation, for us actually to specify them.

This process of derivation and explanation is at the core of our scientific view of the world. For we accept our scientific hypotheses, and so acknowledge the deeper and more basic properties in terms of which these explanations are cast, partly because this enables us to derive and so explain the observable properties we are familiar with. (Indeed our awareness of what we observe is now saturated with a sense of the deeper explanations which pertain to it.) So if conscious events are neural events, then it is natural to suppose that we should also be able to explain their manifest inner features in some such way.

In accepting physicalism, therefore, we apparently incur the obligation of showing how the inner properties of experience can be derived from the physical properties of the neural events we take experiences to be. This is the problem which Nagel expresses by saying that 'If physicalism is to be defended, *the phenomenal features themselves must be given a physical account.*' (Nagel 1974, p. 478) Such an explanation would be a derivation, showing how the phenomenal properties of experience were to be seen as consequences of their physical nature.

(c) The relation between reconciliation and derivation

The problems of reconciliation and derivation are clearly connected. If it seems impossible from the outset to derive the manifest inner properties of experience from physical properties of the brain, this may be because the properties to be derived seem incompatible or incongruent with those from which we aim to derive them. Conversely, if we were to solve this problem

of derivation, then this would perforce also constitute a solution to the problem of reconciliation. For clearly to derive B from A is thereby to show that A and B are fully compatible.

Still the problems are partly distinct. In particular, it may be possible to show how two sets of properties are compatible with one another by some means other than deriving one from the other. (A proof that A and B are consistent is different from a proof that A entails B.) So it might be possible to solve the problem of reconciliation independently of the problem of derivation, and so to treat the two separately. This is how we will approach them in what follows.

From the outset, however, we should note that both these problems seem difficult, and that we have a strong intuition that the problem of derivation cannot possibly be solved. For when we contemplate the physical and chemical properties of brains (or neurons, or whatever) on the one hand, and the manifest inner felt property of pain on the other, we feel that there is no prospect of deriving the latter from the former. How could physics, chemistry, physiology, or the like, provide materials for deducing *how this pain feels (internally, subjectively, privately) to me*? We have no idea what such an explanatory derivation might be like, and we are inclined to feel confident in advance that none could be given.[4]

6. Mental concepts and physical phenomena

Following Davidson many physicalists regard the mental as a conceptual rather than an ontological category. The idea is that we use a distinct and irreducible family of concepts – mental concepts – in thinking of mental phenomena as such, that is, *as mental*. Thus we use the concept pain to think of events *as pains*, the concept belief to think of states *as beliefs*, the concept of desire to think of states *as desires*, the concept of a self to think of things *as selves*, etc. These concepts are interwoven in their use, for when we think of a pain as a pain, we also think of it as something going on in ourselves, etc. The events, states, and other items to which we actually apply these concepts, however, are physical ones, centred in the brain and nervous system. (And generally, since mental and physical phenomena are identified via their causal roles, the physical phenomena to which mental concepts apply encompass the causal roles we assign to the mental) Hence on this account mental concepts actually serve us as *ways of*

[4] Thus as Nagel also says, 'it seems clear *in advance* that no amount of physical information about the spatiotemporal order will entail anything of a subjective, phenomenological character.' (Nagel 2000, 439)

thinking of physical phenomena. They are the mechanisms which evolution has given us for thinking about the causes within us of what we say and do – and hence for thinking about what happens in our own brains – in advance of any explicit knowledge about the insides of our bodies in general, or our nervous systems in particular.[5]

Putting the matter this way makes it clear that we can regard mental concepts both as *irreducible to physical concepts*, and also as *physical concepts (that is, concepts for physical phenomena) of a particular kind*. If we take physical concepts to be those which we use to think of things *as physical* (where we oppose this to thinking of things *as mental*) then mental concepts may be irreducible to physical concepts, in the sense that there are neither definitions nor strict natural laws relating their applications; and the same may hold for the types, properties, or states which we use these concepts to demarcate. A main lesson of philosophy, indeed, is that very many concepts – including everyday concepts like 'clock' or 'cup' – are indefinable or irreducible in this way. Reducibility is a form of strict redundancy, and this seems relatively uncommon among human concepts: like words, particular concepts have particular uses, and the use of one seems rarely to be just the same as another.

So mental concepts can also rightly be regarded as physical concepts (and hence as determining physical types, properties, or states), in the sense that the concepts themselves are physical (neural) forms of representation, and these serve to represent things which are also physical. In this sense mental concepts and mental phenomena are as physical as rocks and stones and trees. They are part of the physical guiding mechanisms of physical living bodies. This is their biological nature as we know them and there is no more reason to hold that they might have an alternative non-biological or non-physical nature than there is to hold that there might be non-biological or non-physical trees. This holds even if the state- or property-types we impose in commonsense psychological thinking are distinct from those of everyday physicality or the physical sciences and irreducible to them. For as non-reductive physicalists stress, questions of conceptual reduction are distinct from questions of ontological nature. The irreducibility of the concept 'cup' to those of chemistry or physics does not tend to show that cups are not physical things.

[5] I have discussed the representation of the mind as inner in relation to evolution in my 2000b.

7. Different kinds of irreducibility

Mental concepts are, however, subject to more than one kind of irreducibility. To take one main kind, applicable to representations generally, consider photographs and the concepts/types we use in describing them. We ordinarily describe photographs *intentionally*, that is, in terms of what they are *about* or *of*. These are the objects and properties represented in the photographs, and (where real objects are concerned) which played the appropriate causal role in their production. We also, however, speak of such photographs as *in the camera*, i.e. as now realized in the photographic apparatus itself, in the form of exposed film, or the registers of digital memory. So in addition to describing a photograph intentionally, we can also describe it intrinsically and physically, in terms of chemistry of the exposed plate or film, the physical structure of the digital memory, and so on.

These two forms of description – in terms of the things and situations represented, and in terms of the physical vehicles via which this is done – apply to representations of many kinds. Simple and straightforward as they are, they are characteristically irreducible, in the sense that they are neither interdefinable nor related by natural law. For we can see that a given range of objects and properties – Aunt Maude swimming at a particular time and place at Brighton Beach for example – might photograph in many different ways (might cause and be registered in very many different sorts of chemical or electronic structure), and again that physically isomorphic photographs might nonetheless be *of* or *about* different things (eg of a wax model of Aunt Maude at Brighton Beach, etc.) So if someone told us that he had discovered a natural law linking photographs of Aunt Maude at Brighton Beach with a particular chemical or electronic structure, we would know in advance both that the claim was wrong and how to set about producing an actual falsification of it. So we should be prepared to acknowledge that there are no strict nomic generalizations – e.g. no photographico-physical laws – linking intentional and physical types.

This kind of irreducibility can be said to hold over representational concepts, types, or properties; but it clearly shows no important gap in our understanding of such representations, nor any serious limitation on our scientific aspirations with regard to them. We understand what the relevant representations are and how they work, and in context we can determine what we want to know about a given representation in a way which is limited mainly by the material complexity of the case. Likewise such

irreducibility gives rise to no problem of representational causality. We understand photographs as mechanisms established by the relevant physical processes, and relate their representational and causal roles (e.g. in driving a printer) as fully as we need be. This is enough for us to regard our practice of describing photographs or other comparable representations (eg those in computers) as admitting an entirely unproblematic integration with science. In these cases we may lack (what philosophers have classically regarded as) *reducing descriptions*, but we can nonetheless frame *realizing descriptions* in a way which renders the lack of reduction scientifically irrelevant.

This is also a main aspect of the irreducibility which Davidson noted as characteristic of mental concepts such as those of *concept*, *intention*, *belief* and *desire*.[6] For concepts, beliefs, and desires are to be understood as involving forms of physical (neural) representation, which we describe in terms of the objects and situations they represent, and which can also be described in terms of their realization in our brains and nervous systems. As with photographs, we can know in advance that no strict laws will relate our commonsense intentional descriptions of these representations with the more intrinsic descriptions which neuroscience may ultimately enable us to frame. And as also with photographs, this kind of irreducibility in itself presents no barrier to the integration of commonsense psychology with neuroscience. But the innerness of the mental relates to an irreducibility of a deeper kind, as we will consider shortly.

8. Conceiving physical innerness as mental

We began with our puzzling commonsense notion of the double innerness of the mental. So let us now note that we can apply the standard physicalistic approach above to this as well. In thinking of the ways mental items are *in* us, we are applying the commonsense mental concept of innerness which we have used from childhood, and which, as Wellman's investigations suggest, may develop naturally in our species. Accordingly we should suppose that this concept has the same general function as our other mental concepts, namely that it serves for thinking about the causal mechanisms within us which regulate our behavior.

[6] I have simplified here by omitting Davidson's early stress on rationality as a source of irreducibility in favour of his later emphasis on intentionality. I have discussed this in relation to Davidson in my 1995.

There is an obvious suggestion as to how it does so. The mental concept of innerness enables us to think about the innerness of this system of causes, which in fact is a kind of physical innerness. So in using our mental concept of innerness, we are thinking of physical innerness, but thinking of it *as mental* rather than *as physical*. Just as we can think of internal neural mechanisms which prompt and guide our behavior as desires, beliefs, or feelings, so we can think of the internality of these mechanisms by thinking of them as *in* the mind.

We are inclined to think of mental events as going on (somehow non-extendedly or non-physically) in the mind, and of the mind or self as (somehow non-extendedly or non-physically) inside the body. In fact, as physicalists acknowledge, mental events go on (extendedly and physically) in the nervous system, and the nervous system is (extendedly and physically) inside the body. So the dual but non-physical innerness which we ascribe to mental events – in the mind, in the body – parallels a dual physical innerness – in the nervous system, in the body – which mental events actually enjoy. This is how it would be if our conception of events as (somehow non-extendedly) in the mind served as a way of thinking of their occurrence in the nervous system, and our conception of the mind as (somehow non-extendedly) in the body served as a way of thinking of the innerness of the nervous system itself. (And the nervous system really is – as Descartes said of his mind or self – 'not in my body as a pilot in his ship' but rather 'very closely joined and, as it were, intermingled with [this body]…' (Med VI, p. 56))[7]

This is an hypothesis about the work done by a particular family of concepts. As our concepts of the mental enable us to think about (physical and neural) inner causes of our behavior, but in commonsense mental terms, so our concepts of the innerness of the mental enable us to think about the (physical and neural) innerness of these causes, and in commonsense mental terms as well.

9. Ratifying and reconciling the oppositions

Once we consider this hypothesis, we can see that it provides a good account of the oppositions which seem at the core of the problem of consciousness. On the one hand it ratifies these oppositions, by showing that they are genuine, and relate to real differences in nature. On the other it reconciles them, by showing that the physical differences to which they

[7] Compare Nagel's view that 'the core of the self – what is essential to my existence – is my functioning brain.' Nagel (1986, p. 40)

relate are ones which we are conceiving in mental terms, and which are ultimately consistent with one another, as can be seen when we reflect on their physical reality. This again fits with the overall idea that mental concepts are ways of thinking about inner causes of behaviour.

(a) Inner and outer as related to the nervous system

We can see this clearly in the case of the most basic opposition, that between the inner and the outer. According to our hypothesis this opposition marks a real distinction, between what goes on within the nervous system and what goes on outside it. But if this is so then it is clear that thinking of events as in the mind is actually consistent with also thinking of them as in the extra-mental physical world.

The first problem of opposition was to see how, if mental events were physical, the same events could be both inner and outer, or both internal and external to the mind. This understanding of the opposition resolves the problem. There is a real opposition between the internal and the external, for events cannot in the relevant sense be both inside and outside the nervous system. But if mental innerness is understood as a way of conceiving occurrence within the nervous system, then mental events can rightly be regarded as both internal and external to the mind. A pain is internal to the mind (internal to the nervous system) of the person who has it, and external to all other minds, and part of the external world in general.

(b) Introspection as internal awareness of the inner

Consider next the opposition between what is (inner and hence) introspectible, and what is (outer and hence) externally perceivable. In this we can contrast the roles of what, following Kant, we can regard as outer and inner sense. Antonio Damasio introduces a neurological distinction related to this by asking the reader to look away from the page to the room in front and then back again.

> As you looked up many neural stations of your visual system, from the retinas to the cerebral cortex, shifted rapidly from making neural mappings of the page, to mapping the room in front of you. But when you returned to the page, these components resumed mapping the page again … resulting in different visual images… . However, several regions in your 'body sensing' brain, which has the job of mapping varied aspects of your body, did not change at all in terms of the kind of object they represent … . The moral of this story is that some parts of the brain are free to roam all over the world and to map whatever sound, shape, smell or texture that the organism's design enables them to map. But some other brain parts – those that represent the organism's own

> structure and internal state – are not free to roam at all. It is reasonable to hypothesize that this is the source of the sense of continuous being that anchors the mental self. (Damasio, 2003a, p. 227)[8]

The parts of the brain which 'are free to roam all over the world and map whatever shape, smell or texture' are those related to the systems of outer sense, including hearing, sight, smell, and touch. Those which map 'the organism's own structure and internal state' can rightly be characterized in terms of inner sense. A. D. Craig has recently described them as constituting an *interoceptive system*, which 'generates a direct thalamocortical representation of the state of the body…crucial for temperature, pain, itch and other somatic feelings. This, as he says, 'constitutes a basis for the subjective evaluation of one's condition, that is, 'how you feel'. (Craig 2002, p. 665-6) Kant's notion of inner sense, like our commonsense conception of introspection, evidently encompasses the operations of this system, as well as the less bodily experiences, also conceived as inner, which result from outer sense.

Awareness of things, whether internal or external, involves applying concepts to them, that is, thinking of them in certain ways. Thus in external perception we are aware of trees as trees, of computers as computers, etc.; and in introspection we are aware of pains as pains, of visual experiences as visual experiences, etc. again. From a physiological point of view, however, the acts of thinking or conceiving in the two cases are strikingly different.

In the case of introspection or inner sense both the mental activity of conceiving the objects of awareness and the activities and structures conceived are internal to the same nervous system. Thus when we think of pains as pains from inside, the activity of conceiving, which seems centred in the cerebral cortex, is directed towards events in the same cortex, i.e. those which originate in the stimulation of nocioceptors. Such thinking is thus not mediated by the working of the external sensory systems, or the concepts applied in their use for outer things. (Likewise in conceiving visual experience as visual experience, the cortical activity of conceiving is directed to events in the visual system of the same cortex.)[9] In external perception, by contrast, the conceiving of the things we are aware of is normally directed to objects which are presented to the outer senses. Such

[8] For further discussion by Damasio see his 2003b, pp. 105ff.

[9] Visual experience, which is not part of interoception as discussed by Craig, requires more detailed treatment than space permits here. This is partly indicated in the discussion of visual experience in my 2002a and b.

thinking therefore is mediated by the concepts we employ in using these senses. The objects of awareness in this case are presented to us as parts of the spatially extended world external to the nervous system (and hence to the mind) at work in conceiving them. Thus the distinction between introspection and external perception, like that between the inner and the outer, reflects both the distinctive neural systems involved, and the boundaries of the nervous system itself.

In this perspective what we call introspection is a basic from of self-awareness, in which the neural structures involved in thinking are applied relatively directly to activities and structures in the very same nervous system. (So in this case 'directly' means without mediation by the activities of perceiving and conceiving involved in the working of the outer senses, which are normally not focussed on the nervous system itself.) It is in this way that our most fundamental sense of self – that of Descartes' *cogito*, with its apparent non-extendedness – is generated.

We noted above that the same event (a pain, say) can be both internal and external to the mind (internal to one nervous system and external to another). Now we can add that the same event (the same pain) can be conceived both internally (from inside) and externally (from outside). Such an event is normally internally conceived by the person who has it, and externally conceived – and so thought of as located in space external to the self – by others. But a person might feel pain while observing the very activity that he was thus internally conceiving in a scanning device, and so at the same time perceiving and conceiving it from outside, and as part of the world external to the self. In approaching the same events via these distinct perceptual and conceptual routes he would of course conceive them very differently; and this is a point on which our discussion will turn.[10]

---

[10] It may be worth making some of the distinctions latent in this more explicit. We can do this by sketching an interpretation which is illustrative if overliteral. Let us use proper names such as 'Tom' to designate persons, capitalized words such as 'PAIN' to indicate the application of the concepts accociated with them, 'P' to designate the neural activity involved in simply having pain, 'IC' to designate conception from the inside, and 'OC' to designate conception from outside. Thus if we take Tom's feeling pain, and take it that he internally conceives this pain as pain, we can designate the situation as:

Tom(IC(P)PAIN)

This in turn corresponds to a particular complex of realizing neural activity, namely that of P together with that involved in the IC of P as Pain. We might, however, wish

(c) Phenomenal as physical but conceived introspectively

Now let us take the difference between the phenomenal qualities of experience and the physical qualities of things external to the mind. When we apply concepts to things we thereby represent them as having properties or qualities. The items to which we apply the same concept are perforce alike, and in thinking of them as alike we thereby think of them as having the same property.

In perceiving (and so conceiving) something we regard as physical, our application of concepts is characteristically mediated by our outer senses, and in particular by sight and touch. These senses are adapted to mapping things in the space exterior to the nervous system, and in using them and the concepts informed by them we represent the affordances of objects in the environment for our navigation and manipulation. In this we

---

to stresss that Tom himself is internally aware of himself as feeling pain, and this would give:

Tom(IC(Tom, P) I FEEL PAIN)

This would explicitly designate a more inclusive pattern of neural activity, involving the IC of the self as I as well, arguably tacit in the previous formulation. By contrast when Fred ascribes pain to Tom (and taking Tom to be aware of his own feeling pain) the situation would be

Fred(OC(Tom(IC(Tom P)I FEEL PAIN) TOM FEELS PAIN

This would correspond to a very different pattern of neural activity, involving an OC of Tom's P and, plausibly, of his IC of himself as well.

In these terms if Tom were both feeling pain and observing a scan of the neural activity involved in this he would be engaged in two sets of conceptualizing activity. The first would be the original

Tom(IC(Tom, P) I FEEL PAIN)

The second might be represented as

Tom(OC(Tom, P) I FEEL PAIN) THAT IS THE NEURAL EVENT OF MY FEELING PAIN

The differences between these two activities of conceiving, according to the argument in the text, owe their deepest differences to the fact that no externally applied concept can have the causal and representational role, and so determine the same perspective on the same event, as one which is applied internally.

see and think of things as having colour; three-dimensional shape (and hence surface texture and interior volume and countless further internal properties ready to hand); location in relation to the perceiver's body; and so on. In perceiving and thinking of things *from outside*, that is, we think of them, and in a particularly full way, as objects extended and arranged in the physical space about us. (As Kant says, the form of outer sense is space.)

All this contrasts with the introspective representation of items internal to the body and nervous system which manifest phenomenal properties. These are hidden beneath the skin, and their role in our lives is not that of things external to the body which it is relevant for us to see or touch. So we (or our brains) do not irrelevantly represent these inner events as visible or tangible, nor do we show their affordances for manipulation or navigation as parts of an extended spatial array. Instead, as in the case of pain, we represent them as manifesting properties which have a unitary and direct influence on behaviour, and which are therefore radically distinct from those which are visible, tactile, or the like. So although neural events of pain do occur in physical space, our interoceptive representations do not show them as spatial (or fully spatial): as having insides or outsides or surfaces or volume or as in a space populated by other extended things. Rather we tend to represent only part of the neural activity connected with the location of some injury, and that location mainly simply as hurting.

This may explain why philosophers such as Nagel and McGinn so emphasize the problem of relating consciousness to space. Nagel urges that mental concepts 'don't give us the comfortable initial handle on the occupants of the familiar spatio-temporal world that prescientific physical substance concepts do',[11] while McGinn stresses the necessity for ' "containing" the non-spatial (as we now conceive it) phenomenon of consciousness.'[12] The idea that phenomena such as pain are not spatial, however, would involve a non-sequitur: a move from the fact that pains (or conscious experiences more generally) are not conceived introspectively as having the spatial features linked with touch and sight, to the idea that they *lack* such features, and are therefore not extended, and so non-physical. Since conscious events are neural events, consciousness is certainly a spatial phenomenon; and this is consistent with the fact that in introspection we do not represent things as having a full complement of spatial properties. All this – the distinct neural and conceptual systems

[11] Nagel (1998, p. 339).
[12] McGinn (1995, p. 159).

involved, their different objects, and the different ways they work – entails that the internal properties we ascribe to experiences such as pain in our own case must perforce be radically different from the externally perceptible and spatially integrated properties of things we perceive as part of the external world. This deep sense of difference is partly registered in our regarding this property as phenomenal as opposed to physical. But of course the neural activity we represent as phenomenal is just as physical as anything else. What we regard as phenomenal is just an aspect of the physical, as we conceive it introspectively. (In this perspective, therefore, we can also regard dualism as mistaking *a partial conceptualization of something which is physical* for *a full conceptualization of something which is non-physical*.)

(d) Subjective as objective but introspectively conceived

Again, consider the difference between the subjective and the objective. When we feel a pain, as we have noted, the pain itself does not seem distinct from our feeling of it, and its whole nature seems disclosed in this feeling. This certainly contrasts with perceiving something external, whose existence is independent of us and our perceiving, and whose nature is only very partially revealed to our senses. But again the difference is best understood physically, and in terms of that between introspection and external perception.

In the external case we naturally make a sharp distinction between our mental activity in conceiving and the thing itself which we conceive. As noted we see this thing as external to us, and as having spatial wholeness, form, and relations, and so a host of other properties bearing on what we can do with the thing or in relation to it. In the case of pain, by contrast, the internal activity of conceiving and the internal event conceived form parts of *a single extended neural event*. We think of this, moreover, mainly in terms of its direct effect on us, as opposed to what this effect tells us we might do via manipulating something else. Hence we think of it as a spatially uncomplex but behaviourally directive impingement of painfulness. So the objective and richly propertied neural occurrences we conceive as pain appear *in our conception* as simple and subjective: their existence seems to us scarcely distinguishable from our feeling of them, and we represent them as having no other nature than we feel them to have.

This explains why pains and other experiences should seem subjective (or why they should seem to have, in John Searle's phrase, a 'first-person ontology'). As before, this may explain our tendency to think of pains,

etc., as subjective *as opposed to* objective, and hence as not physical. But again the conclusion does not follow. Rather the first-person subjectivity of introspected events is a consequence of their objective neural nature; and our sense of fully grasping their internal aspect is consistent with their having further qualities, of which in introspection we are simply unaware.

(e) Private as public but introspectively conceived

Finally, take the distinction between the private and the public. When I feel pain, it seems impossible that anyone else should feel *this* pain, or should be aware of *how this pain feels* in the same way as I am. So the pain seems private, and this seems inconsistent with its being physical and public. This, however, is another reflection of the facts we have already considered. A person who feels pain applies the concept of pain internally – to activity in her own nervous system – and thereby represents the instantiation of what we regard as a phenomenal property. It is in the physical nature of the case that no one else can likewise introspectively apply that concept, and so as to yield that instance of that property. Such first-person conceiving can be done by the subject alone. But this again derives from the neural nature of events like pain, and so is not inconsistent with, but rather explicated by, their physicality. The private is public, but introspectively conceived.[13]

11. Phenomenology as the presentation of causal role within the self

The roles of internal and external perception are linked with a further topic, namely the way phenomenology manifests the working of causality within the self. One of the most striking features of mental concepts is the way they simultaneously face in two directions: inwards towards the self, and outwards towards the other, and hence also towards the world apart

[13] It is perhaps worth noting that this does not make the inner aspect of pain – or the inner application of the concept pain to one's own neural activity – private in the sense opposed by Wittgenstein's arguments about private language. These were directed against a language which was private in the sense that it could not be interpreted by another, or again against the idea of a word denoting something that only one person could apply a concept to. But on the present account events of pain or feeling pain are public physical events, to which more than one person can apply concepts. Insofar as the inner conceptual activity of conceiving and the inner activity of pain itself are distinguishable, then each admits of investigation, as does their combination. So the inner aspect of pain manifest solely in the first-person perspective also admits of public consideration, if not of external conceptualization which would reduplicate the internal. The relation of the present account to Wittgenstein's is briefly discussed in my 2000a cited above.

from the self. This again reflects the difference between applying such a concept to activity in our own nervous system, and applying it to activity in the nervous systems of others, via the way this activity is related to observable behavior and circumstances.

We regard the mental events which go on in us as having both phenomenology and causal role. In thinking of a sensation *as pain*, we think of that sensation both as having a certain internal character, and as thereby affecting the person who has it in ways related to that internal character itself. We take the person who has the pain to conceive the internal character of pain and to react to it *from inside*, whereas we others conceive it (and otherwise react to it, if we do) *from outside*. In this we make use of a single concept – the concept *pain* – which can be applied both from within and without.

The concept which we use in these ways, moreover, is in both uses both a phenomenological and a causal concept. For the phenomenology of mental events – the way they feel to us, the way they are presented in consciousness – has direct causal import. In feeling pain, as noted, we feel our selves *impinged upon*, and in a way particular to the pain we feel; and in this we feel ourselves as immediately engaged in a causal process involving the self. So in feeling pain on touching something we thereby feel ourselves as *moved* to withdraw from the contact, in regarding something as a cause of pain we thereby feel ourselves as *moved* to avoid it, etc. (The matter is likewise, except opposite, when we feel pleasure in contact with something – in this case, by contrast, we feel ourselves moved to continue the contact or to renew it.[14]) The same holds, in a variety of ways, for all the varieties of experience. To experience pain or any other sensation is to respond to the operation of an internal cause.

Thus phenomenology should not be regarded as distinct from causality; rather it is *intrinsically* causal. Phenomenology is the presentation of causes as internal to, and hence as part of, the self – the inner face of causality as it manifests itself from within the self. In encountering such causes from within we thereby gain further purchase on them, and so on our selves. Our conceptual awareness of these inner causes gives us the opportunity to work with them – and so to bring other internal causes to bear on them – in feeling, thought, and will.

[14] We can put this point by saying that phenomenology is what Millikan calls 'pushmi-pullyou' representation: representation whose biological function is to have some effect or other, in this case an effect on the self.

10. The problem of derivation and the world apart from the self

This bears on the second aspect of the problem of consciousness, namely the problem of derivation. This is the problem described by Leibniz in the 17th Century as the impossibility of using features of a mechanism 'to explain a perception', and by a long series of philosophers and scientists since. The kind of explanation apparently required is a derivation of the phenomenal features of experience from physical or physiological features of the nervous system: for example, a derivation of the phenomenal property of pain from the chemical or physical features of neural firings. And as we have already stressed, we have a strong intuition that no such derivation is possible. As Du Bois-Reymond asked in the 19th Century:

> What conceivable connection is there between certain movements of certain atoms in my brain on one side, and on the other the original, indefinable, undeniable facts: 'I feel pain, feel lust; I taste sweetness; smell the scent of roses, hear the sound of an organ, see redness'? (Du Bois-Reymond, p. 45)

We can now see that this lack of 'conceivable connection' is a natural consequence of the way we distinguish the inner and the outer, and hence draw the boundary between the self and the world apart from it. For as we noted at the outset, we apply our notion of mental innerness at first level to the mind or soul, and hence to the self itself. In thinking of ourselves as selves, we regard ourselves as subjects who experience the world and agents who seek to change it, and we regard these phenomena of experience, thought, and will as going on in our selves and bodies. We saw that the subjectivity and privacy we ascribe to inner events – the way their existence and essential nature seems barely distinguishable from our apprehension of them – reflects the way the act of conceiving and the event conceived are so integrated as to constitute an extended event in a single nervous system. This same causal-conceptual integration evidently also integrates the range of neural events of which we are aware in such a way as to present them as aspects or parts of ourselves.

So in thinking of events as inner we thereby also construe them as *our own*, as we do a pain, a visual or auditory experience, or a feeling of fear or pleasure. In thinking of things as outer and physical, by contrast, we think of them as *parts or aspects of the world which is distinct from the self*, and in this we set them apart from the self as internally demarcated. As we saw in considering the phantom limb, the boundary between self and not-self seems to correspond roughly with the skin. But since we draw this boundary from inside, by our inner use of mental concepts, it can seem

intact even in the absence of the body; and the innards of the body themselves can seem external by comparison with it. Hence when we think of our experiences from inside they seem part of ourselves, but when we think of them from outside, as events in our physical nervous systems, they again seem distinct and apart from us.

It is because our distinction between the inner and the outer is also our distinction between the self and the world apart from the self that we feel in advance that the demand for a physical (outer) explanation of the inner qualities of experience cannot be met. The demand actually contains a conceptual or a priori absurdity, which we experience as a dissonance in imagining it fulfilled. For it is in effect a demand that we should somehow replicate the direct and spatially reduced inner mode of cognition by which we conceive aspects of the working of our own nervous systems as parts of our selves, by the external and more fully spatialized mode by which we conceive things as parts of the world distinct from the self. Asked, in effect, to envisage some future thinking of the inner self via the conceptual scheme we use in thinking of the outer non-self, we naturally draw a blank. For since these two ways of thinking actually have contradictory properties – and in particular since one is causally direct and immediate and the other is not – this is a demand which could not possibly be met.

To see this more clearly it may be useful to consider the question in another form. C. D. Broad observed that an angel who knew the whole of mechanical physics would not be able to deduce from it how ammonia smelled, and Frank Jackson has likewise stressed that an omniscient but sensorily deprived neuroscientist would not be able to deduce how pain feels.[15] If Broad's angel, or Jackson's Mary, were in possession of a physical explanation of the phenomenal properties of experience, they could use this to derive the knowledge they lacked, and also, if they wished, to investigate further topics in philosophical phenomenology, such as what it would be like to be a bat. Our conviction that they could not do this is the same as our conviction that there is no such explanation as might make it possible. But what tells us there can be no such derivation is not our knowledge of science and its limits. Rather it is our apprehension that when we feel pain we are in a kind of basic and direct causal-conceptual contact with an aspect of ourselves – a kind of contact realized, as we see, only among parts of a single nervous system; and we rightly feel that this self-involving causal-conceptual directness would be lost in thinking of any different or more scientific kind.

[15] See Broad (1947, 71) and Jackson (1986). I have altered Broad's example to fit the present discussion.

12. Conclusion: reconciliation without reduction or derivation

Examining the innerness of the mental thus enables us reframe the problem of consciousness in two ways. At the outset we distinguished two problems. The first was that of reconciling the inner and apparently non-physical features of mental events with their physical ones, as represented by Nagel's claim that 'We do not at present possess the conceptual equipment to understand how subjective and physical features could both be essential aspects of a single entity or process.' The second was that of deriving the inner and apparently non-physical features from physical ones, as represented by Nagel's claim, carried forward from Leibniz and many others, that 'If physicalism is to be defended, the phenomenal features themselves must be given a physical account.'

If our argument so far has been correct, we should reject both these claims, and reconsider what is problematic about consciousness accordingly. As regards the first, we have seen that if we take our remarkable but commonsense notion of the double innerness of the mental as a way of representing the physical innerness which mental events actually enjoy, then the opposition between the subjective and physical features of conscious events can be both ratified and resolved. The apparently anti-physical subjectivity of these events derives from the direct and spatially reduced way in which we think them as mental, and we accordingly take their inner properties as paradigmatic of their mental nature and essential to it. But since the events and properties we conceive in this way are actually neural and physical, we can rightly take this as essential to their nature as well.

As regards the second claim, we see that in conceiving events as inner we also conceive them as parts or aspects of our selves, as opposed to the world apart from the self. This internal relation to the self is constituted by the same causal and conceptual directness and integration – the same confinement to a single nervous system – as underpins their subjectivity. This mode of conception is unique and part of the particular causal flow which we conceive in using it. The notion that we should somehow replicate it via another – by physical explanation or any form of reduction to the outer – is not a scientific ideal to which we should aspire (or which we might suppose others to attain) but a conceptual incoherence which we should recognize as such. This is not to say that consciousness and its relation to the physical presents no problems. But it suggests the

importance of framing these problems in a way that enables us to recognize more clearly what they are.

References

Augustine: *Confessions.* New York: Sheed and Ward, 1942.

Broad, Charles D. 1947: *The Mind and its Place in Nature*. London: Kegan Paul, Trunch, Trubner & Co.

Chalmers, David 1995: 'The Puzzle of Conscious Experience'. In: 'The Hidden Mind', *Scientific American*, 273, pp. 80-86.

Craig, A. D.: 'How do you feel? Interoception: the sense of the physiological condition of the body'. *Nature Reviews Neuroscience*, 3, pp. 655-666.

Damasio, Anthony R. 2003a: 'The person within'. *Nature*, 423, 15, p. 227.

— *2003b: Looking for Spinoza: Joy, Sorrow, and the Feeling Brain*. New York: Harcourt.

Descartes, *Meditations on First Philosophy: With Selections from the Objections and Replies,* ed. by Cottingham, John. Cambridge: Cambridge University Press 1996.

Du Bois-Reymond, Emile: 'About the Limits of Natural Knowledge' in his *Vorträge über Philosophie und Gesellschaft*, Hamburg: Meiner 1974. Quoted and translated by Bieri, Peter in 'Why is consciousness puzzling'. In: Metzinger, Thomas (ed.), *Conscious Experience*.

Hopkins, Jim 1995: 'Irrationality, Interpretation and Division'. In: MacDonald, Cynthia and MacDonald, Graham (eds.): *Philosophy of Psychology: Debates on Psychological Explanation*. Oxford: Blackwell Publishers.

— 2000a: 'Psychoanalysis, Metaphor, and the Concept of Mind'. In: Levine, M. (ed.), *The Analytic Freud*. London: Routledge, pp. 11-35.

— 2000b: 'Evolution, Consciousness, and the Internality of Mind'. In: Carruthers, Peter and Chamberlin, Andrew (eds.): *Evolution and the Human Mind*. Cambridge: Cambridge University Press, pp. 276-299.

Jackson, Frank 1986: 'What Mary Didn't Know'. *The Journal of Philosophy*, 83, pp. 291-95

McGinn, Colin 1995: 'Consciousness and Space'. In: Metzinger, Thomas (ed.): *Conscious Experience*, pp. 149-163.

Metzinger, Thomas 1995: 'The Problem of Consciousness'. In: Metzinger, Thomas (ed.), *Conscious Experience*. Schoningh/Imprint Academic, pp. 3-37.

Nagel, Thomas 1974: 'What is it like to be a bat?' reprinted in his *Mortal Questions*, Cambridge: Cambridge University Press, 1979, pp. 165-180.

— 1986: *The View from Nowhere*. New York and Oxford: Oxford University Press.

— 1998: 'Conceiving the Impossible and the Mind-Body Problem'. *Philosophy*, 73, pp. 337-352.

— 2000: 'The Psychophysical Nexus'. In Boghossian, Paul, and Peacocke, Christopher (eds.): *New Essays on the A priori*. Oxford: Oxford University Press, pp. 433-478.

Ramachandran, Vilayanur S. and Blakeslee, Sandra 1998: *Phantoms in the Brain: probing the Mysteries of the Human Mind*. New York: William Morrow Publishers.

Reid, T. 1785: *Essays on the Intellectual Powers of Man,* ed. by Brookes, Derek R. University Park, Pennsylvania State University Press 2002.

Wellman, Henry 1993: 'Early Understanding of the Mind: the Normal Case'. In: Baron-Cohen, Simon, Tager-Flusberger, Helen, and Cohen, Donald J. (eds.) *Understanding Other Minds, Perspectives from Autism.* Oxford: Oxford University Press, pp. 10-39.