Journal of Consciousness Studies
www.imprint-academic.com/jcs

Ulrich Mohrhoff

The Physics of Interactionism

Physics has been invoked both to refute and to support psycho-physical interactionism, the view that mind and matter are two mutually irreducible, interacting domains. Thus it has been held against interactionism that it implies violations of the laws of physics, notably the law of energy conservation. I examine the meaning of conservation laws in physics and show that in fact no valid argument against the interactionist theory can be drawn from them. In defence of interactionism it has been argued that mind can act on matter through an apparent loophole in physical determinism, without violating physical laws. I show that this argument is equally fallacious. This leads to the conclusion that the indeterminism of quantum mechanics cannot be the physical correlate of free will; if there is a causally efficacious non-material mind, then the behaviour of matter cannot be fully governed by physical laws. I show that the best (if not the only) way of formulating departures from the 'normal', physically determined behaviour of matter is in terms of modifications of the electromagnetic interactions between particles. I also show that mental states and events are non-spatial, and that departures from the 'normal' behaviour of matter, when caused by mental events, are not amenable to mathematical description.

I: Introduction

There is another hard problem, in addition to the problem of how anything material can have the subjective, first-person phenomenology of consciousness (Chalmers, 1995). It is the problem of how anything material can have freedom. By 'freedom' I mean a person's ability to behave in a purposive, non-random fashion that is not determined by neurophysiological structure and physical law. I do not mean the absence of other determining factors, as this would render freedom synonymous with randomness.[1]

I decide to raise my right arm and up it goes. This decision — a mental event — appears to me to be both the cause of the ensuing physical event and a causal primary. (A causal primary is an event the occurrence of which is not necessitated by antecedent causes.) I can think of various reasons for raising my arm (I may want to catch a ball), and these may involve antecedent causes (e.g., the ball was thrown in my direction); but if there is anything that made it inevitable that I should raise my arm, I know nothing of it.

[1] I have no quarrel with compatibilism, the view that free will is compatible with determinism. My freedom may well consist in being governed by what I intrinsically am (what the Indian contemporaries of Plato and Aristotle would have called my 'self-nature' or 'self-law', *svabhava*, *svadharma*) rather than by universal laws or a combination of universal laws and randomness.

Journal of Consciousness Studies, **6**, No. 8–9, 1999, pp. 165–84

To be sure, ignorance of an antecedent cause does not prove its nonexistence. But what does? We can aspire to establish that events of type *C* are regularly followed by events of type *E*. If we succeed, we are free (are we?) to imagine a hidden string between individual events such that each event of type *C* is the cause of an event of type *E*. Failure to establish the existence of a type of event *C* the instances of which are regularly followed by events of type *E*, on the other hand, is not a proof that events of type *E* lack antecedent causes. It doesn't prove anything beyond our ignorance of the antecedent causes of events of type *E*. Proving an event a causal primary is an impossible task. But from this it does not follow that there are no such events. What does follow is that empirical science cannot aspire to know that an event is a causal primary. That is why scientists may ignore causal primaries. But this is no reason for philosophers to dismiss their possible existence.

The absence of causal primaries is often called 'the causal closure of the physical world', where 'physical' means 'non-mental' rather than 'governed by the laws of physics'. This causal closure is a trivial consequence of the lack of scientific interest which results from being unable to identify causal primaries. What is causally closed is the scientifically known world, not the world as such. Yet there are many philosophers who look upon causal closure as an ontological truth and go on to invoke it as an argument against interactionism, the doctrine that mind and matter are two mutually irreducible, interacting spheres. Interactionists as I understand them are motivated primarily by a desire to make room for free will, the denial of which is both counter-intuitive and at odds with notions (moral and otherwise) that are central to the fabric of our active lives. While repudiating the causal closure of the physical world, interactionists nevertheless shrink from contesting the validity of the laws of physics, not realizing that this is contingent on the presumption of causal closure.

That is how we come to witness futile fights between philosophers of mind who reject interactionism on the ground that it is incompatible with the laws of physics and, in particular, the law of the conservation of energy, and interactionists who meekly defend their position, claiming that, by exploiting the loophole of quantum-mechanical indeterminism, non-material mind is capable of influencing matter without violating conservation laws. Both the charge and the defence are misconceived. The law of energy conservation is either true by virtue of the meaning of 'energy', and therefore is not threatened by interactionism, or it is contingent upon the causal closure of the physical world, and therefore is no threat to interactionism. The loophole hypothesis, on the other hand, violates basic physical laws other than the law of energy conservation.

This should not come as a surprise. To be causally efficacious, mental events that are causal primaries must make a difference to the behaviour of matter and thus to the behaviour of its constituent particles. The effects of such events on the behaviour of particles have to be expressed in the language of physics, for this is the only language suitable for describing the behaviour of particles. But the laws of physics presuppose causal closure and describe the behaviour of matter in the absence of causal primaries. Hence it follows that the behaviour of matter in the presence of a causally efficacious non-material mind cannot be fully governed by those laws.

The hard problem of consciousness and the hard problem of freedom appear at first sight to be logically independent. To embrace the irreducibility of consciousness, one need not deny the causal closure of the physical world, and one need not attribute to consciousness a causal role, as has been stressed by Chalmers (1997). On the other hand, it is possible to have physical events interspersed with non-physical causal primaries that lack subjective properties. Take Eccles' (1994) theory in which 'psychons' in the mind affect physical processes in the brain. As Chalmers (1997) has pointed out, the question of whether psychons have any experiential qualities is irrelevant to the causal story.

But this apparent independence of causality and subjectivity is called into question every time someone utters the word 'consciousness'. To see this, suppose that consciousness is irrelevant to the causal story. Then it is explanatorily irrelevant to our claims about consciousness: the physical act of making a judgement about experience is not sensitive to the experience itself (Kirk, 1996). In other words, there are two mutually irrelevant kinds of experience, the experience^E that we actually have and the experience^L about which we make statements. Zombie philosophers make judgements about experience^L, but it would be self-contradictory for them to conceive of the distinction between experience^E and experience^L. These conclusions seem to constitute a *reductio ad absurdum* of the supposition that consciousness is irrelevant to the causal story. Hollywood, it seems, has got it right: zombies are shuffling affectless brutes, not smart philosophers of mind (DeLancey, 1996). Taking the hard problem of consciousness seriously thus appears to make it necessary to take the hard problem of freedom as seriously.

And so we have more than sufficient reason to address the latter problem as vigorously as the former has been addressed in recent issues of this journal and elsewhere. In the present article I will apply myself to the preliminary task of 'deconstructing' physics-based arguments purporting to prove the nonexistence of freedom. Section II reviews the argument from energy conservation — the claim that it is inconsistent with the interactionist doctrine — and the counter-argument that purports to show that interactionism and free will are consistent with the unbroken reign of physical law.

Section III refutes the arguments against interactionism that invoke conservation laws. It begins with an examination of what physicists mean by 'energy' and 'momentum'. The respective conservation laws are shown to be consequences of these meanings. They are necessarily true whenever 'energy' and 'momentum' are well-defined concepts. For these concepts to be well defined, it is however not necessary that the quantities they denote are conserved everywhere and under all circumstances. If they fail to be conserved, it can be for either of two reasons. It may be that energy and momentum are indeed meaningless; the curved space–time of Einstein's general theory of relativity provides an instructive example of this possibility. Or it may be that they are conserved somewhere but not everywhere. Then they are meaningful even where they are not conserved, as for example where matter is causally open to a nonmaterial mind.

Section IV refutes the argument purporting to show that quantum mechanics offers a way of reconciling interactionism with the unbroken reign of physical law. According to this argument, an intention to act can be causally efficacious by merely modifying the probabilities associated with individual quantum events. I show that, on the contrary, an intention to act cannot be causally efficacious without modifying the statistics of ensembles of such events. And this is the same as saying that it cannot be causally efficacious without modifying some physical laws.

Section V shows how the departures from the laws of physics due to non-material mind can (and must) be formulated in the language of physics. The appropriate mathematical entity is the electromagnetic four-vector potential (or, simply, the electromagnetic field). As a summary representation of possible effects on moving particles that makes no reference whatever to causes, the electromagnetic field necessarily represents the effects of both material and non-material causes. Section VI shows that a causally efficacious non-material mind is not something that exists in space, and that its action on matter is not amenable to mathematical description.

A more technical discussion of the physics is available (Mohrhoff, 1997).

II: Energy Conservation and the Interactionist Hypothesis

Attempts to address the mind–body problem along interactionist lines have traditionally been faulted for taking liberties with physical conservation laws, notably the principle of the conservation of energy (also known as the first law of thermodynamics). M. Bunge (1980, p. 17) and D.C. Dennett (1991, p. 35) speak for the prosecution.

> If immaterial mind could move matter, then it would create energy; and if matter were to act on immaterial mind, then energy would disappear. In either case energy . . . would fail to be conserved. And so physics, chemistry, biology, and economics would collapse.

> Let us concentrate on the returned signals, the directives from mind to brain. These, *ex hypothesi*, are not physical; they are not light waves or sound waves or cosmic rays or streams of subatomic particles. No physical energy or mass is associated with them. How, then, do they get to make a difference to what happens in the brain cells they must affect, if the mind is to have any influence on the body? A fundamental principle of physics is that any change in the trajectory of any physical entity is an acceleration requiring the expenditure of energy, and where is this energy to come from? It is this principle of the conservation of energy that accounts for the physical impossibility of 'perpetual motion machines', and the same principle is apparently violated by dualism. This confrontation between quite standard physics and dualism has been endlessly discussed since Descartes' own day, and is widely regarded as the inescapable and fatal flaw of dualism.

Dualists have taken these strictures to heart. Even Karl Popper, by proclaiming himself not to be 'in the least impressed by the danger of falling foul of the first law of thermodynamics' (Popper and Eccles, 1983, p. 564), implicitly acknowledges the danger. From the early days of quantum mechanics, the strategy of the defence has consisted in claiming that quantum-mechanical indeterminism allows non-material mental events to act on matter (specifically the brain) without violating conservation laws. Eddington (1935) was probably the first to speculate publicly that the mind may influence the body by affecting quantum events within the brain through a causal influence on the probability of their occurrence.

More recently H. Margenau (1984) has suggested that the mind may be 'regarded as a field in the accepted physical sense of the term', yet not be 'required to contain energy in order to account for all known phenomena in which mind interacts with brain' (p. 97): 'In very complicated physical systems such as the brain, the neurons and the sense organs, whose constituents are small enough to be governed by probabilistic quantum laws, the physical organ is always poised for a multitude of possible changes, each with a definite probability' (p. 96).

Standard axiomatizations of quantum mechanics recognize two kinds of change: the probabilistic collapse of a quantum-mechanical superposition which occurs during a measurement, and the deterministic evolution of the quantum state which takes place between measurements (von Neumann, 1955). Margenau proposes that the causal efficacy of mind rests on the following sequence of steps: (i) The relevant physical system develops, in accordance with the deterministic evolution of states, into a superposition of alternative states, each associated with a probability. (ii) Mind alters the physically determined probabilities, possibly by superimposing its own probability field on the physically determined probability field. (iii) The resulting superposition collapses to one of its elements in accordance with the probabilistic change of states. In this way, Margenau argues, mind can act on the brain without disturbing the balance of energy. D. Hodgson (1996) likewise invokes the mind's ability to load the quantum dice.

Seizing on Margenau's proposal, J.C. Eccles, in collaboration with F. Beck (Beck and Eccles, 1992; Eccles, 1994), has put forward one of the most elaborate and specific hypotheses of mind–brain interaction to date. It capitalizes on the basic unitary activity of the cerebral cortex, exocytosis. Exocytosis is the emission of chemical transmitters into the synaptic cleft by a vesicle of the presynaptic vesicular grid, a paracrystalline structure situated inside the terminal expansion (bouton) of a nerve fibre. It is an all-or-nothing event, which has been found to occur with a probability of about one fourth to one third when a bouton is activated by a nerve impulse. Eccles and Beck assume this probability to be of quantum-mechanical origin. They cite increasing evidence for a trigger mechanism that may involve quantum transitions between metastable molecular states, and propose a model for the trigger mechanism based on the tunnelling of a quasi-particle through a potential barrier.[2] According to their model, during a period of the order of femtoseconds the quasi-particle is distributed over both sides of the barrier. One side corresponds to the activated state of the trigger, the other side to the non-activated state. At the end of this period exocytosis has been triggered with the aforesaid probability. Eccles and Beck propose that mental intentions act through a quantum probability field altering the probability of exocytosis during this brief period.

While the postsynaptic effect due to the change in probability of exocytosis by a single vesicle is many orders of magnitude too small for modifying the patterns of neuronal activity even in small areas of the brain, there are many thousands of vesicles per bouton and many thousands of similar boutons on a pyramidal cell (the principal type of neuron of the cerebral cortex), and there are about 200 neurons in the region of a dendron, the basic anatomical unit of the cerebral cortex (Eccles, 1994, p. 98). The hypothesis of mind–brain interaction according to Eccles and Beck is that mental intention becomes neurally effective by momentarily increasing the probabilities for exocytosis in the hundreds of thousands of boutons in a whole dendron.

> In summary it can be stated that it is sufficient for the dualist-interactionist hypothesis to be able to account for the ability of a non-material mental event to effect a changed probability of the vesicular emission from a single bouton on a cortical pyramidal cell. If that

---

[2] The microtubule hypothesis adopted by Penrose (1994) is a membrane-physiological proposal for this trigger mechanism, as F. Beck (1994) has pointed out. It realizes the motion of the quasi-particle as the motion of one, or a few, hydrogen atoms in the membrane.

> can occur for one, it could occur for a multitude of the boutons on that neuron, and all else follows in accord with the neuroscience of motor control (Eccles, 1994, p. 78).

> It is reassuring that all of the richness and enjoyment of our experiences can now be accepted without any qualms of conscience that we may be infringing conservation laws! (Eccles, 1994, p. 170).

III: Conservation of Energy and Momentum: A Closer Look

Originally, momentum was defined as 'mass-times-velocity'. It soon became apparent that (within Newtonian physics) this was a conserved quantity. Then the special theory of relativity superseded Newtonian physics, and mass-times-velocity was no longer conserved. By this time, however, the property of being conserved was accorded much greater importance than the original definition in terms of mass and velocity. Momentum accordingly was redefined so as to match its original definition in the low-speed limit, where the two theories make identical predictions, as well as to retain its status of a conserved quantity.

But a redefinition that consists in the substitution of one theory-dependent *definiens* for another, can only be a halfway stop. It must be possible to define the *definiendum* at a more basic level, independently of the specific principles of either theory and hence in a way that is valid for both. It indeed soon transpired that the different mathematical embodiments of momentum in the respective theories of Newton and Einstein were specific instances of a quantity that could be invariantly defined for a large class of theories. In 1918 E. Noether discovered a deep connection between symmetries[3] and conservation laws. This exists in all theories that can be derived from a mathematical expression known to physicists as the Lagrangian. In all such theories (and these include not just all experimentally well-confirmed theories to date but all theories esteemed worthy of consideration by contemporary physicists), a continuous symmetry[4] implies the existence of a locally conserved quantity.[5] And one of these locally conserved quantities implied by the continuous symmetries of the Lagrangian is called 'momentum'. Thereafter it was possible to claim that this has always been the true definition, even when the concept was insufficiently differentiated from its then sole instantiation, mass-times-velocity.

The same holds true of energy. Both energy and momentum are defined as conserved quantities. They are conserved by definition. Either they make sense and are conserved, or they don't make sense. They don't make sense whenever the mathematically described world (or, equivalently, the Lagrangian) does not possess the symmetries that imply their respective conservation laws; in other words, whenever the corresponding symmetry transformations, applied to a mathematical description of a physical situation, yields not just a different description but a different physical situation.

---

[3] In physics a symmetry is both a consequence and an expression of the fact that the mathematical description of the world is underdetermined by observational data. Just as a symmetrical figure can be transformed into itself (for instance, by a rotation), so a mathematical description of the world can be transformed into a different mathematical description of the same world (for instance, by a rotation of the coordinate system).

[4] 'Continuous' here means that the corresponding transformation, like a rotation of the coordinate axes, can be carried out continuously rather than in discrete steps only.

[5] Saying that a physical quantity is locally conserved is the same as saying that the amount of it inside any bounded region of space can change (from B_1 to B_2, say) only if the difference B_2–B_1 passes through the boundary of the region.

The symmetry that gives meaning to 'momentum' is known as the homogeneity of space; it consists in the mechanical equivalence of all locations in space, or in the fact that every closed mechanical system behaves in the same way anywhere. The symmetry that gives meaning to 'energy' is known as the homogeneity of time; this consists in the mechanical equivalence of all moments of time, or in the fact that every such system behaves in the same way anytime. Translate the coordinate origin in space and/or time, and what you get is a different description of the same physical situation. This has the nature of a postulate: differences in the outcomes of identical experiments performed at distinct locations and/or times are to be ascribed to the different physical conditions (known or unknown) prevailing at these distinct locations and/or times, not to these locations and/or times *per se*. An instance of the synthetic *a priori* judgement that everything that happens has a cause, this postulate has more to do with what we (investigating humans) make of our experiences than with any particular experience of ours. If we did not assume the existence of a cause, we would not look for one; and if we did not assume the existence of physical causes to explain the spatial or temporal inhomogeneities we observe, we would not look for such causes but rest content with attributing those inhomogeneities to space or time *per se*.

And so it would seem that the homogeneity of space and the homogeneity of time are *a priori* certain; that momentum and energy are therefore always well defined; and that they are always conserved. However, there are riders to this series of conclusions. Whatever is *a priori* certain is so only with regard to our mental constructs. Whether or not these can be thought of as descriptions of objective reality is another matter. Also, before anything can be derived from the said homogeneities, they must be given formal expression within the framework of a physical theory. And there is no *a priori* guarantee that this is possible. In fact, there are reasons to surmize the opposite, as will become apparent in what follows.

There is nothing controversial about the way in which space and time are rendered manifestly homogeneous (that is, the way in which their homogeneities find mathematical expression in a physical theory). Either one introduces a privileged class of coordinate systems (called 'inertial systems') or one lets a mathematical entity known as the metric tensor (or simply, the metric) do the privileging (by taking a particularly simple form in the privileged systems). However, what is capable of manifesting homogeneity also lends itself to the manifestation of inhomogeneities. The metric needed to manifest the flatness[6] of space or space–time could instead serve to manifest the curvature of a Riemannian space or space–time. This is the same as saying that the metric texture of space or space–time offers a handle for the formulation of an interaction law. Matter could act on matter via the intermediate representation of the metric in much the same way as electric charges act on electric charges via the intermediate representation of the electromagnetic field. The curvature at any space–time point *p*, determining partly if not fully the motion of matter at *p*, could depend on the distribution and motion of matter elsewhere and at earlier times. It could thus represent a causal influence on the motion of matter at *p* due to the earlier distribution and motion of matter elsewhere.

---

[6] It will become evident below that the 'flatness' of space(–time) is actually the same as the mechanical equivalence of locations heretofore called the 'homogeneity' of space(–time).

This *a priori* possibility is an actual feature of the objective world. The interaction in question is gravity; the theory just outlined is the general theory of relativity. Now, gravity appears to be quite indispensable to the creation of what Squires (1981) has called an 'interesting world'. Without gravity there would exist no stars, no planets, nor (for all we can imagine) any sites hospitable to something as interesting as life. In view of this it might be asserted that curvature is implied by our own existence, or that since we are here, space–time cannot be flat.

At any rate, the metric connection lends itself to the manifestation either of spatio–temporal homogeneity or of gravity. As far as the description of objective reality is concerned, the choice is not ours but Nature's. And Nature has opted for gravity. The metric which could have offered a handle for the incorporation, in our mental picture of reality, of a homogeneous space and a homogeneous time, is already used up. From this and what has been said earlier one might draw the conclusion that in situations in which gravity plays a significant role, energy and momentum are undefined. But such a conclusion would ignore that even curved space–time is locally flat,[7] and that, as a consequence, the energy and the momentum of all non-gravitational fields are locally (as opposed to globally) conserved. This is sufficient for them to be well-defined. What is ill-defined in any generic space–time is the gravitational energy/momentum, and hence the total energy/momentum. The energy/momentum associated with a curved region of space–time is, strictly speaking, definable only in model space–times that are flat 'at the edges'.[8]

At certain junctures in the history of physics the law of energy conservation has been called in question. Bohr at one time felt that he had to renounce it, and not a few particle physicists despaired of it before the neutrino was proposed and, in due course, discovered. It should not be supposed that these physicists were unaware of the deep connection between the conservation laws for energy and momentum and the homogeneity of time and of space. Rather they were driven to consider the possibility that these homogeneities were not, after all, respected by Nature. Bohr thought that the problems facing atomic theory were 'of such a nature that they hardly allow us to hope that we shall be able, within the world of the atom, to carry through a description in space and time that corresponds to our ordinary sensory perceptions' (in Honner, 1982). If the feasibility of such a description cannot be taken for granted, the homogeneity of space and of time cannot be taken for granted either.

More recently, in connection with the so-called measurement problem in quantum mechanics, the stochastic generation (and hence non-conservation) of energy has emerged as a theoretical possibility (Ghirardi *et al.*, 1986; Pearle, 1989). This amounts to introducing stochastic inhomogeneities in the 'flow' of time, and to redefining energy as the quantity whose conservation would be implied if those inhomogeneities were absent. If such a definition is adopted, the view that the conservation of energy is part of the meaning of 'energy', can no longer be entertained.

The situation, then, is this: If the energy conservation law is part of the meaning of 'energy', the interactionist hypothesis cannot imply a violation of this law. And if

---

[7] Reduced to two dimensions this means that any sufficiently small (infinitesimal) patch of a smoothly curved surface is approximately (exactly) flat.

[8] This does not mean that gravitational energy/momentum cannot be approximately defined whenever and wherever space–time can be considered as approximately flat. If it can, the departures from flatness can be treated as a gravitational field in flat space–time.

physicists can invoke inhomogeneities in the 'flow' of time and define energy in such a way that it is conserved only when and where those inhomogeneities are absent, interactionists can do the same. The causal efficacy of non-material mind could be based on its generating similar (but not stochastic) inhomogeneities. As long as there exists an experimental realm in which mind-generated inhomogeneities are absent or negligible (and from a physicist's point of view, given present experimental limitations, they may well be negligible everywhere), energy remains well-defined even where matter is causally open to non-material mind. If no such realm existed, attributing energy to matter would be gratuitous, since in this case any mathematical expression would do. None could be tested, because the proof that one has the right expression lies in the experimental corroboration of its conservation. But if the formula for the energy associated with matter is testable somewhere, nothing prevents one from using the same formula everywhere, including where matter is open to the action of non-material mind and energy is not necessarily conserved.

IV: Interactionism Violates Physical Laws

While the argument from energy conservation does not succeed, the notion that mental events can influence physical events through the loophole of quantum-mechanical indeterminism, without in any manner whatsoever infringing on the deterministic regime of physical laws, is chimerical, as is shown presently.

Consider a causally efficacious mental event (say, the intention to flex the right index finger). If this occurs in the mind associated with any healthy body, the intended action takes place. If the same intention occurs in the minds associated with an ensemble of healthy bodies, all of those bodies flex their right index fingers as a result. There is no randomness in the causal concatenation between intention and intended action. Throughout the ensemble, the same mental event brings about the same physical event.

Consequently, if the causal efficacy of a mental intention is postulated to involve modifications of quantum-mechanical probabilities associated with 'collapsible' wave functions, these modifications are statistically significant. In the simplest case in which the modifications amount to the selection of one out of two possible outcomes in a single collapse, the same outcome is selected every time the intention occurs. In the Eccles–Beck model, in which the intended action is the effect of many weak modifications, accumulated over a large number of collapses, the fact that the same action is produced every time entails that the individual modifications likewise exhibit statistically significant trends.

A clear distinction must be maintained between sets of active sites in the same brain and the statistical ensembles of active sites relevant to the present discussion. The latter involve different brains or, more precisely, different instances of identical brains. Consider an ensemble of such brains. Then consider an ensemble of vesicles such that each vesicle is from a different brain and all vesicles occupy identical positions in their respective brains. There are as many such vesicle ensembles as there are vesicles in each brain. Let us compare two cases. In the first case all brains are influenced by a certain mental intention; in the second case none of the brains is influenced by it, other things being equal. What needs to be compared in particular is the behaviour of each vesicle ensemble in the two cases. If none of the vesicle ensembles shows any difference in the percentage of 'firing' vesicles, the intention cannot be causally efficacious. If it is causally efficacious, the intended effect takes place whenever the intention is present in the minds associated with those brains, and only then. In this case there must be some vesicle ensembles for which the percentages of 'firing' vesicles differ in the two cases.

In a word, if single-case probabilities get modified, there are statistical ensembles whose behaviours get modified. What gets modified is not merely individual quantum events but the statistics of entire ensembles of such events. And these statistics, unlike the individual events, are fully determined by physical laws. Changing them means changing the physical laws.[9] Altering the single-case probabilities associated with individual measurement-like events without changing the laws of physics is possible only if the relative frequencies associated with every ensemble of identical such events remain unaltered. But this is possible only if the individual modifications of probability are themselves probabilistic. Suppose that some of the single-case probabilities are increased and some are decreased such that the overall probability remains unchanged. Then the laws of physics remain unchanged, but there can be no talk about causation, mental or otherwise. Whatever 'causes' such statistically insignificant modifications of probability cannot be causally efficacious. To be causally efficacious, an event must make a difference every time it occurs. It must make a difference to the behaviour of some ensemble, that is, it must be statistically significant. The basic tenet of the interactionist position — causal openness of the material to the non-material mental — thus entails a violation (that is, an occurrence of modifications) of physical laws.[10] Probability distributions, determined jointly by initial conditions and some quantum-mechanical equation of motion such as the Schrödinger equation, are altered. One might leave it at that. But one might also wonder if any such alteration could not be formulated just as well in terms of the well-known physical quantities that determine probability distributions during the deterministic phase of their evolution. This is the case, as I proceed to show.

V: Interactionism without Quantum Collapses

As an illustration of how the altered probability distributions entailed by the interactionist hypothesis could arise within the formalism that physicists use to calculate probability distributions, rather than as *ad hoc* modifications of the results of the calculations, we will now consider an open one-particle system. A system consisting of just one particle obviously cannot accommodate the creation or annihilation of particle pairs, but it seems reasonable to assume that minds do not cause either type of event. (The energy needed for pair creation is available in cosmic rays and

[9] Measurements on ensembles of identical quantum-mechanical systems evolving under identical initial and boundary conditions yield identical distributions of results. Modified statistical distributions observed on ensembles of identically prepared systems indicate modified boundary conditions. Modified boundary conditions can arise from modifications either of the spatial distribution of environmental matter or of the fields generated by this matter. Modified boundary conditions given identical such distributions (that is, modified fields) imply a modification of the physical laws according to which the fields are generated.

[10] The same violation is entailed when the non-material self is replaced by a 'superintelligence' who, as F. Hoyle (1983) surmizes, guides the evolution of the cosmos by altering the probabilities associated with quantum processes.

high-energy physics laboratories, not in brains. The antiparticles needed for annihilation events are not normally present in brains.) We further assume that mental events do not induce particles to change type. This is tantamount to ruling out the so-called strong and weak forces as vehicles of mental causation, for it is these that cause type conversions. (The weak force can for instance convert electrons into neutrinos.)

The strong and weak forces are unlikely vehicles of mental causation because both of them are short-range forces. The strong force is confined to the interior of certain subatomic particles, the mesons and the baryons. A residue of this force, the so-called nuclear force, is confined (in brains if not in neutron stars) to the interior of the atomic nucleus, as is the weak force. None of these forces is effective at the scale of chemical processes; none therefore is relevant to the chemistry of the brain. The goings-on inside atomic nuclei have no influence on when neurons fire, or how likely they are to fire, which is how the causal efficacy of the mind must make itself felt.

And since the most general formulation of effects on the motion of a spinless particle already includes the possible effects on a particle with spin, we can confine our discussion to that type of particle which is represented by a single wave function (rather than one of those multicomponent wave functions known as spinors). Such a particle is known as a scalar particle.

The entire physics of a quantum-mechanical system is formally contained in a mathematical expression known as the probability amplitude. This amplitude allows physicists to calculate (at least in principle) the likelihood with which the system transits from any initial state to any final state in any given interval of time. The entire physics of a scalar particle is in fact known if one knows the amplitude associated with how likely the particle is to travel from point x to point y in any given time span.

It is a remarkable fact about quantum mechanics that this amplitude (let's represent it by the symbol $<y|x>$) can be calculated by 'summing over' (that is, adding up contributions from) all space–time curves that connect x at the starting time with y at the time of the particle's arrival — as if the particle went from x to y by travelling along every possible path (Feynman and Hibbs, 1965). Each curve simply contributes a complex number of unit magnitude. Such a number is fully specified by what is called its phase. The phase of a curve is the sum of the phases associated with its segments, and this fact makes it possible to think of the phase of a curve as its **length**. For an uncharged particle this **mechanical** length of a curve in space–time is simply proportional to the **geometric** length of the same curve, and the proportionality factor is simply the particle's mass.[11]

Clearly, the only way of influencing the motion of a scalar particle (charged or uncharged) is to modify the mechanical lengths of curves in space–time.[12] This can be done in one of two ways: in the manner of gravity, by changing the geometric lengths of curves and thereby warping space–time itself, or by changing the mechanical lengths without changing the geometric lengths.[13] When it is weak

[11] As is customary among theoretical physicists, we pretend that some universal constants are equal to 1.

[12] This is the reason why homogeneity (or the mechanical equivalence of locations in space–time) is tantamount to the flatness of space–time.

[13] Admittedly it is difficult for non-mathematicians to see how the same curve can have different lengths, a geometric and a mechanical one, and how it can even have different mechanical lengths for different

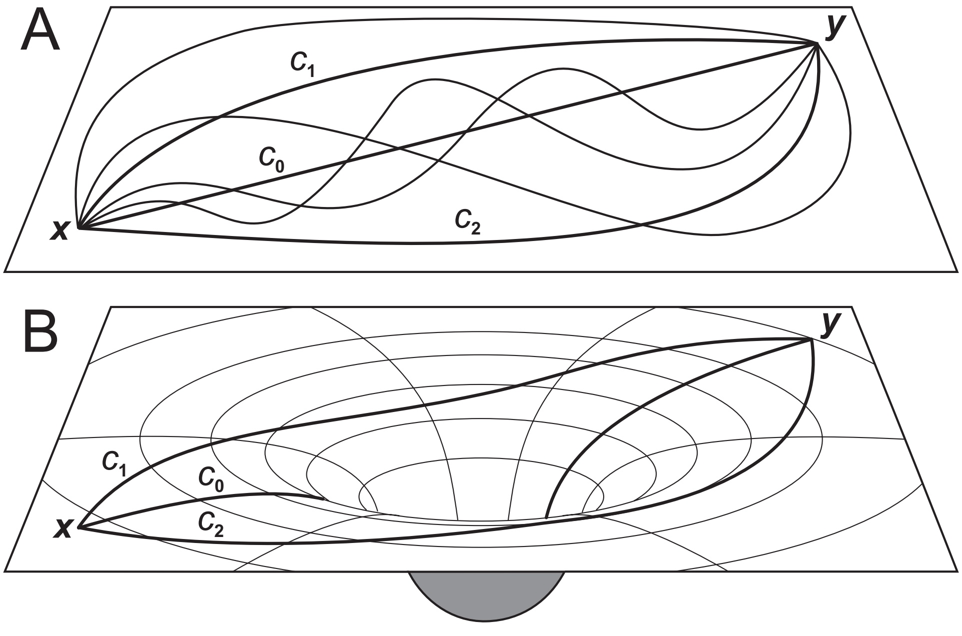

Figure 1.
The upper diagram (A) shows a few of the curves contributing to <*y*|*x*>, the amplitude associated with the probability that a particle initially located at *x* is later found at *y*. By Euclidean standards, the shortest curve is the straight line c_0. The possible effects on the motion of a particle are mathematically represented by a non-Euclidean way of measuring lengths. In terms of *mechanical* lengths, the shortest curve connecting *x* and *y* may be c_1, or it may be c_2, or it may be c_1 for particles of a certain type and c_2 for particles of a different type.

Diagram B: Because gravity affects all particles alike, its effect on the mechanical lengths of curves can be thought of as a warping of space-time itself. The surface with the dip represents space-time. The extra dimension into which it is warped is not physical; its sole purpose is to make it possible to visualise the warping of space–time. The dip could be due to a massive object at its centre. Because of the dip, c_0 is no longer the shortest curve connecting *x* and *y*. A classical particle travelling from *x* to *y* will take the shortest curve on either side of the dip, and this makes it seem as if a force, gravity, were pulling the particle towards the centre of the dip as it travels around it.

enough to permit a human brain to function normally, gravity plays no significant role in a region of space the size of a brain, which is why we only need to consider the latter option.[14]

---

types of particle. As a useful analogy, consider all the routes from Zurich to Copenhagen, say. There are (at least) three ways of measuring their 'lengths': in kilometers, in hours, and in litres of petrol. One route may be the shortest in the ordinary sense, another route may be the shortest as measured in hours, and yet another may be the shortest in terms of petrol consumption. It is even more difficult to see how a particle can behave as if it travelled simultaneously along all possible curves connecting two locations. It should however be born in mind: (i) that this is a rather straightforward description of one of the most successful mathematical formalisms used in physics; (ii) that there is no reason whatever to expect visualisable models drawn from everyday experience to be appropriate for dealing with the extreme limits of human experience; and (iii) that a consistent realistic interpretation of the formalism is yet to be found, if one can be found at all.

[14] Further reasons for dismissing gravity as a possible vehicle of mental causation are given below.

As an illustration of the kind of effect caused by changes in the mechanical lengths of space–time curves, imagine a plane (see fig. 1A). In it imagine two points x and y and a bundle of curves beginning at x and ending at y. One of these curves (call it c_1) will have a shorter mechanical length than every other curve. By no means does this have to be the straight line c_0. Next suppose that the mechanical lengths of all curves are increased in such a way that those of curves entirely to the left of c_1 increase more than those of curves entirely to the right of c_1. As a result, the mechanically shortest curve will no longer be c_1 but a different curve c_2 to the right of c_1. One of the effects of altering the mechanical lengths of space–time curves is thus equivalent to bending the curve of minimum mechanical length between any two space–time points. (Usually there is just one such curve, but see fig. 1B for a situation in which the shortest curve connecting x and y is not unique.)

In the so-called classical limit, in which quantum mechanics degenerates into classical mechanics, the only contributions to $<y|x>$ that 'survive' come from the curve (or curves) of minimum mechanical length. (More precisely, from curves that are shorter than their nearest neighbours.) This explains why a classical particle travels from x to y (in the specified time span) along the mechanically shortest curve (or one of the mechanically shortest curves) between x and y. What gets bent are the space–time trajectories of classical particles. But bending the space–time trajectory of a classical particle is the same as accelerating the particle, and this is the reason why in classical physics one talks about acceleration-causing forces instead of modifications of mechanical lengths.

How does one mathematically represent modifications of the mechanical lengths of curves that leave the geometric lengths unchanged? The answer is straightforward: by means of some field.[15] This field (let's call it A) associates with every infinitesimal curve segment (depending on both the location and the direction of the segment) the extra bit of mechanical length that the segment has for a charged particle. (For an uncharged particle, recall, the mechanical length of the segment is simply its geometric length times the particle's mass. Uncharged particles do not 'experience' the non-gravitational modifications of mechanical lengths.)

The field A is known to physicists as the electromagnetic vector (or four-vector) potential. It contains exactly the same information as the electric and magnetic fields together.[16] The electric field is what bends the projections of classical trajectories on space–time planes that include a time axis (that is, it accelerates charges in a fixed direction), while the magnetic field is what bends the projections of classical trajectories on spatial planes (that is, it accelerates charges in directions perpendicular to their directions of motion).

The vector potential (equivalent to the electromagnetic field) is thus the summary representation of all possible non-gravitational effects on the motion of a scalar particle, including all effects caused by mental events. Physicists habitually associate the

---

[15] Reminder: mathematical details can be found in a companion article to this paper (Mohrhoff, 1997).

[16] A pedant would mention that the experimental phenomenology is in one-to-one correspondence with the electric and magnetic fields (different fields giving rise to different observable effects) but not with the vector potential. The latter has extra degrees of freedom due to that symmetry of the Lagrangian that implies the conservation of electric charge. The origin of those extra degrees of freedom is readily seen: if the mechanical lengths of all curves from x to y are changed by the same amount, the curve of minimum mechanical length — and with it the physics — remains unchanged.

vector potential not only with the way in which it influences the motion of charged particles but also with a particular way (given by Maxwell's laws) in which it is generated by the motion and distribution of charges. They don't question (and as physicists, concerned solely with the behaviour of inanimate matter, need not question) the assumption that this is also the only way of generating it. But, in fact, anything — be it physical, mental or whatever — that has a (non-gravitational) effect on the motion of a particle, necessarily contributes to the electromagnetic vector potential.[17] If a mental event is to influence the behaviour of the quasi-particle in Eccles' model of a trigger mechanism for exocytosis, it must modify the barrier — a potential barrier — penetrated by the quasi-particle.

When the electromagnetic field was introduced by Maxwell, it was thought of as the property of a mechanical substrate pervading space. When Einstein discarded this substrate, the erstwhile property became a physical entity in its own right. The symbol took on a life of its own; the mathematical description took the place of the thing described. Today many physicists believe that reality is mathematical. While the present investigation ought not to be biased in favour of any such metaphysical claim, it is safe to say that the empirical reality investigated by science is, first of all, a complex of mental constructs. (I am not saying that it is 'nothing but' mental constructs.) What these constructs have in common, and what distinguishes them from mere fantasies, is that they are objectifiable, that is, they are capable of being thought of as features of an objective world. The vector potential is such a construct (after quantization, at any rate), and from the role it plays in our account of particle motion it is clear that it cannot be partial to any particular type of causal agent. It serves to represent the effects of mental causes just as well as those of physical causes.

Now that we know that the second manner of modifying the mechanical lengths of space–time curves is, in actual fact, the way of the electromagnetic force, we have another reason for dismissing gravity (the first manner) as irrelevant to mental causation. Considering that exocytosis is controlled by the influx of Ca^{2+} ions into a synaptic vesicle (Eccles, 1994, pp. 149–53), mental causation is likely to be effected through a modification of the physically determined forces exerted on ions (that is, on charges), particularly those involved in the propagation of nerve impulses. But the electromagnetic interaction between, say, two protons is about 10^{36} times stronger than their gravitational interaction. Hence if the mentally generated modification of the force exerted on a charged particle were of gravitational nature, the mental self would have to generate an implausibly strong gravitational field (about that many times stronger than the physically generated one), while it would only need to generate an electromagnetic field that is weak in comparison with the physically generated one.

Yet another reason why the electromagnetic interaction is the more likely vehicle of mental causation is the selectivity of the electromagnetic force. While this acts on charges only, gravity affects everything. If one wants to make an ion move

---

[17] D. Papineau (1996) writes: 'The central problem facing any contemporary dualist is that twentieth-century science denies any causal powers to unreduced phenomenal properties. Phenomenal properties differ in this respect from electromagnetic forces.' As a matter of fact, they don't. The effects of irreducible phenomenal properties on scalar particles are included in the electromagnetic field.

through a neutral medium, one had better not also accelerate the medium, as this would simply cause a congestion; if one tries to move both the ion and the medium, nothing will move.

However, all said, nothing fundamental stands in the way of the notion that the mind contributes to any or all of the four fundamental forces, inasmuch as the weak and strong forces no less than the metric tensor and the electromagnetic vector potential are simply ways of formulating possible effects on the behaviour of particles, whether their origin be physical, mental or whatnot. For reasons indicated above I believe however that the electromagnetic field is the single most effective vehicle of mental causation, and that therefore the other possibilities are not worth considering.

If non-physical causes do indeed contribute to the vector potential, the well-known dynamical laws of the vector potential (that is, Maxwell's laws or their quantum-mechanical counterparts) are violated, in the sense that they describe some but not all contributions to the vector potential. It is worth emphasising that there are neither theoretical nor experimental reasons to rule out such a violation. While empirical evidence of non-physical contributions to the vector potential may as yet be lacking, absence of evidence is not the same as evidence of absence. Evidence of absence is not available because systems in which such contributions might occur are notoriously complex, difficult to analyse, and no less difficult to experiment with. It could be argued, moreover, that if the non-physical contributions to A amounted to a substantial modification of the physically determined component of A, mind would be able to actuate matter through a less complex physiology. While the complexity of the body is no argument against interactionism, it certainly suggests that a non-material mind cannot cause more than minute modifications of the physically determined component of A.

As for theoretical derivations of the dynamical law for A, they tell us no more than what was initially assumed. Because A can be considered as a quantum-mechanical system in its own right, its dynamics is known if one knows how to calculate the amplitude for the transition from any initial field configuration to any final field configuration in any given time span. As there are contributions to the amplitude $\langle y|x\rangle$ from all curves connecting x and y, so there are contributions to this transition amplitude from all 'histories' of the field A (that is, from all curves in the infinite-dimensional space of field configurations). And as before, each contribution only depends on the mechanical length of the corresponding history/curve.

A crucial difference however arises when it comes to finding the correct mathematical expression for the mechanical lengths of field histories. The formula for the mechanical lengths of space–time curves 'experienced' by a scalar particle contains the representation of all possible effects on the motion of a scalar particle. We can be sure that none have been left out. On the other hand, we can be sure that we have the right formula for the mechanical lengths of field histories only if all sources contributing to the field are represented in it, and only if the effects represented by the field are linked to their causes according to universal mathematical laws. In order to be able to derive Maxwell's equations (along with their quantum-mechanical counterparts) we must therefore assume (i) that the motion of a particle cannot be affected by anything except the motion and distribution of particles, and (ii) that the action of particles on particles is amenable to mathematical description. Hence the argument that mind cannot affect the behaviour of charged particles because this is governed by Maxwell's laws, obviously begs the question.

VI: Mind, Space and Mathematical Description

The aim of this section is to show (i) that mind is non-spatial and (ii) that the action of a non-material mind on matter is not amenable to mathematical description. The latter conclusion in fact is a consequence of the former.

On the interactionist view, mind is non-spatial and, as causal agent, independent. ('Independent' here means that its acts of will are not fully determined by physiological microstructure and physical law. 'Non-spatial' means that the mental cause of an effect on the motion of particles in the brain does not consist in the spatial distribution and/or the state of motion of objects in space.) From this it follows that the condition that the effects represented by the electromagnetic field must be linked to their causes via universal mathematical laws, cannot be satisfied for the direct effects of volitions. For one thing, if this condition is to be satisfied, the causes must be, at the very least, amenable to mathematical description. Since this is essentially synonymous with spatio–temporal description, they must have positions in space. For another thing, if the link between a causal primary in the mind and its physical effect were amenable to mathematical description, one could write down a Lagrangian for the mind as causal agent. But if that were possible, this causal agent would be just another kind of matter subject to just another kind of physical law — something whose existence neither dualists nor materialists are likely to endorse. For the dualists, it would be too materialistic; for the materialists, too dualistic.

It is in fact unnecessary to assume the non-spatiality of the mental, as is shown presently. If the self were an object in space, it would have to make sense to talk about the position of the self relative to other objects in space. Let us see why it does make sense to talk about the relative positions of particles. Like the non-material self, a fundamental particle can't be seen. Its position relative to other material objects can nevertheless be inferred from its observable effects, for instance from a trail of droplets in a cloud chamber. But this inference is possible only because there exists a physical law that relates the position of the particle to the positions of its effects. Applying our knowledge of this law to observational data (the positions of the droplets), we can infer the (approximate) position of the particle. And how did we come to know this law? It is an extrapolation from regularities observed in the relative positions of larger charged objects that can be seen.

By the same token, attributing to the self a position appears to make sense only if there exists a law relating the position of the self to the positions of observable effects caused by the self. If we knew such a law, we could infer the self's position from its effects. But how could we discover such a law? By observing regularities in the positions, relative to observable material objects, of larger selves that can be seen? There may be psychophysical laws (Chalmers, 1995) relating mental states to physical configurations in the brain, but so far nobody has suggested that these laws involve the positions of mental states. I suppose that this is because there simply is no way of making sense of the position of a mental state. Only the physical effects that the self, *ex hypothesi*, is capable of producing, are localizable in space.

Not all theorists of consciousness would agree. M. Lockwood (1989, p. 101), for one, takes special relativity to imply that mental states must be in space given that they are in time. This conclusion, however, appears to rest on a too naive identification of two distinct concepts of time. What 'time' means in the context of psychological experience is not the same as what it means in the context of special relativity. Without an in-depth study of their relation (not offered by Lockwood), only the physical effects of mental states can be said to necessarily exist in space–time. See Clarke (1995) for a refutation of Lockwood's arguments in support of the spatio– temporal localization of mental events.[18]

A priori, the modifications of the electromagnetic field 'experienced' by certain constituents of the body could be effected in two ways: the non-material self could contribute to the electromagnetic field as a separate source, or it could modify the way in which the field is built up by material sources. However, to act as a separate source, the self would have to exist in space, and this notion has just been rejected. Hence it follows that material particles are the only sources of the electromagnetic field, and that the non-material self can only influence the summary effect — represented by the electromagnetic field — of the action of particles on particles.

The causal efficacy of the self thus rests on the causal efficacy of the particles, or on the ability of the particles to modify their individual contributions to the electromagnetic field. The causal behaviour of particles (meaning, the way particles influence each other's motion, as distinct from the way particles move) accordingly comes in two modes: a physical mode which obeys the laws of physics, and a non-physical mode through which modifications of the physical mode are effected. But this means that the only causal agents in existence are the fundamental particles, and that the non-material self cannot be as non-material as dualists would have it. Interactionism thus cannot be the last word. The implications of this, as well as the possible relationship between the self and the body's constituent particles, will be explored in another article (Mohrhoff, submitted).

VII: Summary and Outlook

The following results have been obtained:

(1) The conservation of energy and momentum is a consequence of the homogeneity of time and of space. This is warranted for systems that are causally closed. As to material systems that are open to causal influences from non-material mind, either energy/momentum is/are ill-defined or there is no reason why it/they should be conserved.

---

[18] If one thinks of mental representations as non-physical properties of conscious organisms, one may posit a separate non-physical substance as the substrate for those properties. But there is no need to distinguish between a physical and a non-physical substance; the same substance can have both physical and non-physical determinations. It can also have both spatial and non-spatial determinations, for it is only the spatial determinations that are necessarily in space, not their intrinsically indeterminate substrate. Nor (since not even the substrate for physical determinations necessarily exists in space) would a separate substrate for non-physical determinations necessarily exist in space. (Even the fundamental constituents of matter do not necessarily exist in space; there is nothing in the theories or the phenomenology of physics that would contradict the view that space contains the relative positions of fundamental particles rather than the particles themselves.)

(2) Assuming that part but not all of matter is causally open to non-material mind, it makes sense to attribute (non-conserved) energy and momentum even to physical systems that interact with non-material mind.

(3) The causal efficacy of non-material mind implies departures from the statistical laws of quantum physics. These departures are capable of being formulated in terms of modifications, by the conscious self, of the electromagnetic interactions between particles; and they are more consistently formulated in this manner.

(4) Because the electromagnetic field is a summary representation of effects on the motion of particles, the effects caused by mental events are necessarily among the effects represented by it. It is not that one cannot formulate the effects of the self in terms of a separate probability field, to use Margenau's (1984) term. The point is that this field would be indistinguishable from a contribution to the electromagnetic field, which makes it obvious that departures from the laws of physics are involved. Thinking of the effects of the self as contributions to the electromagnetic field is preferable for two reasons. First, it eschews the contentious notion that measurement-like events take placein the unobserved brain. Second, it leads to a more unified treatment of causality. There is no reason whatever for having probabilities determined twice over, once during their deterministic evolution by the physically determined vector potential, and once at the end through a superimposed probability field generated by the self.

(5) Quantum-mechanical indeterminism cannot be the physical correlate of free will. Free will implies departures from the laws of physics.

(6) Mind is non-spatial. There is no point in attributing positions to mental states and events.

(7) The departures from the physical laws caused by non-physical mental events are not amenable to mathematical description. It is worth emphasizing that they are not therefore random. They could be necessitated by something of a primarily qualitative nature, something that manifests itself in quantitative, spatio–temporal terms but is not reducible to these terms.

Although there are no compelling theoretical or experimental reasons why mental events should not be capable of causing departures from physical laws, it may remain difficult for interactionists and proponents of free will, at least for some time to come, to disabuse the contemporary physicist, biologist, or philosopher of science of the doctrine of physicalism, which has been a reigning orthodoxy for well over a century. So much was this doctrine taken for granted, that until recently it was considered as almost indecorous to waste much thought over the dismissal of its antithesis. Thus, after stating that 'very few people any longer suppose that living things violate any laws of physics (as some thinkers supposed as late as the nineteenth century)',[19] Hilary Putnam (1992, p. 83) makes known why this should be so: 'Physics can, in principle, predict the probability with which a human body will follow any given trajec-

[19] Among those 'very few people' are the biologists A. Szent-Györgi (1961), W. Elsasser (1966), M. Delbrück (1986), and Mae-Wan Ho (1993). Their work supports the view that the behaviour of particles in living systems differs from the behaviour of the same particles in inanimate objects.

tory.' Are we to suppose that the mountaineer who fell to his death would have been able to choose a less ruinous trajectory if only Eccles' hypothesis of mind–brain interaction had been true?

What interactionists and proponents of free will claim, in effect, is that the nonmaterial self becomes materially effective by modifying the electromagnetic interactions between constituents of the body. Not only is this consistent with the assumption that the trajectory of the body's centre of mass is fully determined by physical laws, but also it agrees with our sense of free will which interactionists wish to take seriously. I decide to raise my hand and it goes up; but nothing in my experience leads me to expect that I could alter my trajectory once I have jumped off a cliff.[20]

Yet there is cause for optimism. If the hard problem of consciousness is taken as seriously as it now is, the hard problem of freedom is bound to follow suit. Many researchers in cognitive studies now admit the irreducibility of consciousness. And most of the philosophers who speculate about the shape of a fundamental theory of consciousness invoke some form of panpsychism.[21] Yet, with few exceptions, these philosophers still find it necessary to reduce conscious events to 'causal danglers': they affirm that pain is not reducible to its physical correlate yet deny that it causes us to pull our hands out of fires. Such a position is inherently unstable, as Lowe (1995) has pointed out. It is under intense pressure either to lapse back into materialism (which restores the causal efficacy of conscious feelings by identifying them with their physical correlates) or to take the further step of admitting the causal efficacy of consciousness. The present article has shown that, from the point of view of physics, nothing stands in the way of taking this long overdue step.

Acknowledgement

I wish to thank Jean Burns for many helpful suggestions.

References

Beck, F. (1994), 'Quantum mechanics and consciousness', *Journal of Consciousness Studies*, **1** (2), pp. 253–5.

Beck, F. and Eccles, J.C. (1992), 'Quantum aspects of brain activity and the role of consciousness', *Proc. Natl. Acad. Sci. U.S.A.*, **89**, p. 11357.

Bunge, M. (1980), *The Mind-Body Problem* (Oxford: Pergamon Press).

Chalmers, D.J. (1995), 'Facing up to the problem of consciousness', *Journal of Consciousness Studies*, **2** (3), pp. 200–19.

Chalmers, D.J. (1997), 'Moving forward on the problem of consciousness', *Journal of Consciousness Studies*, **4** (1), pp. 3–46.

Clarke, C.J.S. (1995), 'The nonlocality of mind', *Journal of Consciousness Studies*, **2** (3), pp. 231–40.

DeLancey, C. (1996), 'Emotion and the function of consciousness', *Journal of Consciousness Studies*, **3** (5-6), pp. 492–9.

Delbrück, M. (1986), *Mind from Matter? An Essay on Evolutionary Epistemology* (Palo Alto, CA: Blackwell Scientific).

Dennett, D.C. (1991), *Consciousness Explained* (London: Allen Lane / Penguin).

---

[20] My self has something to do with the particles constituting my body, not with particles in other material aggregates. There is thus a clear case for (i) the distinction between interactions that take place between the particles in my body and interactions that involve particles outside my body, and (ii) the supposition that my self is capable of modifying only the former.

[21] For explicit panpsychist proposals see, e.g., Hut and Shepard (1996), Rosenberg (1996) and Seager (1995).

Eccles, J.C. (1994), *How the Self Controls Its Brain* (Berlin: Springer).
Eddington, A. (1935), *The Nature of the Physical World* (London: Dent).
Elsasser, W. (1966), *Atoms and Organism* (Princeton: Princeton University Press).
Feynman, R.P. and Hibbs, A.R. (1965), *Quantum Mechanics and Path Integrals* (New York: McGraw-Hill).
Ghirardi, G.C., Rimini, A. & Weber, T. (1986), 'Unified dynamics for microscopic and macroscopic systems', *Physical Review D*, **34**, pp. 470–91.
Ho, Mae-Wan (1993), *The Rainbow and the Worm: the Physics of Organisms* (Singapore: World Scientific).
Hodgson, D. (1996), 'The easy problems ain't so easy', *Journal of Consciousness Studies*, **3** (1), pp. 69–75.
Honner, J. (1982), 'The transcendental philosophy of Niels Bohr', *Stud. Hist. Phil. Sci.*, **13** (1), pp. 1–29 (Letter to Høffding, 22.9.1923, Bohr Scientific Correspondence, microfilm No. 3, p. 5).
Hoyle, F. (1983), *The Intelligent Universe* (London: Michael Joseph).
Hut, P. and Shepard, R. (1996), 'Turning the hard problem upside down and sideways', *Journal of Consciousness Studies*, **3** (4), pp. 313–29.
Kirk, R. (1996), 'Review of Chalmers: *The Conscious Mind: In Search of a Fundamental Theory*', *Journal of Consciousness Studies*, **3** (5–6), pp. 522–3.
Lockwood, M. (1989), *Mind, Brain and the Quantum* (Oxford: Basil Blackwell).
Lowe, E.J. (1995), 'There are no easy problems of consciousness', *Journal of Consciousness Studies*, **2** (3), pp. 266–71.
Margenau, H. (1984), *The Miracle of Existence* (Woodbridge, CT: Ox Bow).
Mohrhoff, U. (1997), 'Interactionism, energy conservation, and the violation of physical laws', *Physics Essays*, **10** (4), pp. 651–65.
Mohrhoff, U. (Submitted), 'Consciousness and particles: A Vedantic synthesis'.
von Neumann, J. (1955), *Mathematical Foundations of Quantum Mechanics* (Princeton: Princeton University Press).
Papineau, D. (1996), 'A universe of zombies?,' *Times Literary Supplement*, **4864** (June 21, 1996), pp. 3–4.
Pearle, P. (1989), *Physical Review A*, **39**, p. 2277.
Penrose, R. (1994), 'Mechanisms, microtubules and the mind', *Journal of Consciousness Studies*, **1** (2), pp. 241–9.
Popper, K.R. and Eccles, J.C. (1983), *The Self and Its Brain* (London: Routledge and Kegan Paul).
Putnam, H. (1992), *Renewing Philosophy* (Cambridge, MA / London: Harvard University Press).
Rosenberg, G.H. (1996), 'Rethinking nature: a hard problem within the hard problem', *Journal of Consciousness Studies*, **3** (1), pp. 76–88.
Seager, W. (1995), 'Consciousness, information and panpsychism', *Journal of Consciousness Studies*, **2** (3), pp. 272–88.
Squires, E.J. (1981), 'Do we live in the simplest possible interesting world?', *Eur. J. Phys.*, **2**, pp. 55–7.
Szent-Györgi, A. (1961) in *Light and Life*, ed. W.D. McElroy and B. Glass (Baltimore: John Hopkins Press).