The Journal of Mind and Behavior
Summer 1992, Volume 13, Number 3
Pages 211–232
ISSN 0271–0137

How to be a Scientifically Respectable "Property-Dualist"

Ran Lahav

Southern Methodist University

and

Niall Shanks

East Tennessee State University

We argue that the so-called "property-dualist" theory of consciousness is consistent both with current neurobiological data and with modern theories of physics. The hypothesis that phenomenal properties are global properties that are irreducible to microphysical properties, whose role is to integrate information across large portions of the brain, is consistent with current neurobiological knowledge. These properties can exercise their integration function through action on microscopic structures in the neuron without violating the laws of quantum mechanics. Although we offer no positive argument for the existence of irreducibly global properties, the conclusion is that this view is a scientifically respectable hypothesis that deserves to be investigated.

It is often thought that the so-called "dualist" theory of consciousness, according to which conscious experience is a "non-physical" phenomenon, is inconsistent with contemporary science, both in spirit and in detail. In this paper we will argue against this contention. We will attempt to outline certain constraints which, if followed by the dualist, are likely to produce theories that are consistent both with contemporary neurobiological knowledge about the brain and with current theories of physics.

The theory whose scientific respectability we will examine will be what is often termed "property-dualism": the view which construes conscious experience in terms of non-physical *properties* of the brain. We will not deal with

The first part of this paper was written mainly by Ran Lahav, and the second part mainly by Niall Shanks. Requests for reprints should be sent to Ran Lahav, Philosophy Department, Southern Methodist University, Dallas, Texas 75275, or Niall Shanks, Department of Philosophy, East Tennessee State University, Johnson City, Tennessee 37614.

substance-dualism (Cartesian dualism): the view which places conscious experience in non-physical *entities* that are associated with the brain (e.g., the soul). The former seems to us especially close to the spirit of modern science, for reasons that will become clear below. Furthermore, we will concentrate on the *interactionist* version of property-dualism, according to which conscious experiences (construed in terms of "non-physical" properties) causally interact with, and hence exert influence on, their neuronal substratum. We will thus not consider the *epiphenomenalist* version, according to which conscious experience is a by-product of, and is incapable of causally influencing, neuronal events. The reason for this choice is, first, the familiar argument (which will not be discussed here) that epiphenomenalism makes it a mystery how consciousness has developed through evolution if it makes no difference to brain activity. Second, interactionism is commonly considered a more serious violation of modern science than epiphenomenalism, and its defense is therefore more challenging.

Hence, we will talk of consciousness in terms of properties of neuronal entities,[1] and will use a common terminology to call them *phenomenal properties*. Phenomenal properties are those features commonly characterized as what "appears from a subjective point of view," or what "appears in experience" or is "given in consciousness," whether veridically or not. The exact boundaries of the phenomenal domain are not important here, but roughly, they include, presumably, visual shapes, colors, textures and surfaces, auditorily experienced sound-qualities, sensations of pain, feelings of fear, and so on.

The hypothesis we will examine is that phenomenal properties are distinct from the properties posited by standard modern scientific theories. At the same time, they interact with their neuronal substrate. We will suggest that this view is scientifically defensible if it construes phenomenal properties as what we will call *irreducibly global* properties (or, for short, "global" properties). Intuitively, these are properties that apply to global neuronal structures, and exert on them global causal effects which cannot be broken down into more local or microscopic causal effects. The causal effects of such properties are, so to speak, different from the sum of the causal effects of more local properties: the laws that govern the behavior of neurons in systems with such global properties cannot be reduced to, and are not even determined by, the laws that govern isolated neurons. More accurately, using a common philosophical jargon, the causal effects of an irreducibly global property on a neuronal entity do not logically or nomologically "supervene" on (i.e., roughly, are not fixed by) the causal effects of more local properties

[1]Our proposal is neutral between the hypothesis that phenomenal properties are properties of global neural *events* or of global neural *structures* (entities). For the sake of simplicity we will talk only about the latter version.

on this entity.[2] Thus, in a system with irreducibly global properties, the behavior of a neuron depends on the global state of the entire system; not merely on the local conditions in and immediately around it, but also on remote parts of the system.

It is important to note that thus characterized, irreducibly global properties are different from familiar high-level properties, such as macroscopic, organizational, or functional properties of neuronal systems (e.g., the property of being an and-gate neuronal circuit). Such high-level properties of systems make no difference to the individual neuron's behavior—as long as the neuron's internal states and immediate environment are fixed. A neuron does not care, so to speak, about the organizational or macroscopic or functional properties of the system in which it is embedded; it "sees" only its immediate environment. According to the present proposal, on the other hand, a neuron is influenced not only by its neighbors, but also by the global properties of the system in which it is found. Two neurons that are identical in their internal state and immediate environment may nevertheless behave differently if embedded in systems with different irreducibly global properties. Consequently, in such a system, the behavior of the whole is, as it were, different from the sum of the behaviors of the isolated parts.

In the traditional terminology, irreducible phenomenal or experiential properties are sometimes called "non-physical." Theories that assume the existence of such properties are called "property-dualism." However, we find the terms "dualism" and "non-physical" objectionable. If physical events have irreducibly global properties, then there is no reason why they should not be regarded as part of the physical world. After all, irreducibly global properties need not be eerie in any sense that disqualifies them from counting as physical. For this reason we prefer the term "globalism" to "non-physicalism" or "dualism." For the same reason we prefer to call the rival traditional position not "physicalism" or "materialism," but rather "localism." Intuitively, localism is the view that once you fix the properties of the particulars that exist in the universe (as opposed to properties of *sets of* particulars) and their causal powers, you have thereby fixed the behavior of the universe. More accurately, localism is the theory that the causal powers of every property that applies to a complex of entities supervenes (at least in virtue of the laws of nature) on the causal powers of the properties of the component entities of that complex. This excludes, of course, irreducibly global properties.

[2]Roughly, A-properties can be said to supervene on B-properties if any difference with respect to A implies a difference with respect to B, but not necessarily vice versa. Thus, conscious experience is commonly thought to supervene on neural properties, since, presumably, any difference in conscious experience implies a difference in its neural substratum; although the converse does not hold: small changes in neural states need not be reflected in conscious experience. For more details on supervenience see Kim, 1982.

The behavior of a system with irreducibly global properties is determined not only by its local ("physicalist") properties, but also by its global properties.

It is important to emphasize what the paper is *not* designed to achieve. First, we will not attempt to provide any positive argument for globalism ("property dualism") and against localism ("physicalism"). Traditionally, "non-physicalist" philosophers have offered various arguments to show that there is an unbridgable gap between the objective physical world and the subjective, qualitative, or intentional nature of conscious experience. For example, it has been argued (Nagel, 1974; Jackson, 1982) in the so-called argument from knowledge, that one may have a complete knowledge of the physical aspects of an organism, without knowing what its conscious experience is like; which presumably suggests that consciousness is not exhausted by physical phenomena. However, since our point here is the scientific respectability of this view and not its truth, we will not go into these arguments. We hope to provide a logical (or, rather, scientific) space for the traditional "anti-physicalist" intuitions, but how the globalist might try to justify filling this space is another issue which will not be discussed.

Second, our aim is not to provide a full-blown globalist theory, but only some general theoretical constraints within which such a theory might be developed. Thus, many crucial details will remain for the globalist to work out. In particular, although we will show that the globalist is free (within certain constraints) to construe phenomenal properties as having characteristics which localist ("physicalist") phenomena do not have, this leaves it open as to what these characteristics need to be in order to account for their subjective or experiential nature. After all, the mere fact that a property is irreducibly global does not automatically make it experiential or subjective. The globalist needs not only to posit irreducibly global properties, but also to identify them as phenomenal properties. But what this identification amounts to is an issue we will not discuss.

Our discussion will be divided into two interrelated parts: first, the consistency of globalism with neurobiology, and second, its consistency with physics.

Part 1: Globalism and Neurobiology

One might reject off-hand the globalist approach on the grounds that current neurobiological theories leave no room for causal intervention by global properties. Brain functions are controlled by well-known microphysical factors, and there is simply no causal space for any additional controlling factor.

This objection is based on a much too optimistic conception of our current neurobiological knowledge. Although it is true that we have acquired a tremendous number of details about the brain, our knowledge of virtually every neural event—from the behavior of chemical substances to global patterns of neural activity—is limited to no more than general outlines. We have no more than a rough idea of, for example, the spectrum of factors controlling the activity of a single neuron, or of the patterns—and even the general principles—of connectivity of neuronal networks. From the perspective of current neurobiological knowledge, there is more than enough room for global properties to exert causal influences.

However, although the objection does not refute the globalist view, it nevertheless points to a general constraint on any plausible model of global-neuronal interaction. It suggests that the *direct* causal effect of global properties on neural activity is likely to be relatively small; although it may later be amplified into a significant effect. First, since at least the major outlines of known brain processes can be explained within a reasonable approximation in terms of purely chemical and physiological mechanisms, global properties have only a limited leeway to exercise their direct influence. Second, if global properties had complete control over brain activity then the exact structure of neuronal networks, neurons, and sub-neuronal mechanisms would not matter for proper brain-function. In contrast, although our knowledge of the brain is rather scanty, we know that its organization follows quite specific patterns and principles. Assuming that this is not an evolutionary coincidence, and considering the additional fact that subtle damage to the brain may result in severe functional and experiential deficits, it is reasonable that neural organization does much of the work in controlling brain functions. It is likely that global properties—if they exist—are limited to small modulations of neural activity.

This suggests that the globalist should maintain that the function of global properties is limited to fine-tuning neuronal activity. Furthermore, in order to explain why the brain has developed through evolution to have phenomenal properties, it should be assumed that their effect is evolutionarily advantageous for the organism. Thus, the globalist faces two main challenges: first, to suggest a possible neural mechanism through which small influences can be amplified and fine-tune neuronal activity; and second, to suggest a plausible role for such influences. We believe that these challenges can in principle be met within the constraints of present neurobiological knowledge. To show this, some plausible speculations will be offered below.

Phenomenal Properties Apply to Global Systems

There is strong evidence that phenomenal events are associated not with any particular neural structure, but rather with complex patterns of neural activities across diverse parts of the brain. The reason is that even a simple phenomenal field contains a variety of features whose representations are distributed across a variety of neural areas. This can be illustrated through the example of the visual system (the following information is reviewed in DeYoe and Van Essen, 1988; Lahav, 1990, Zeki and Shipp, 1988).

In primates (like in many other mammals), the primary visual pathway, from the retinae to visual and other structures in the cortex, passes through more than twenty distinct visual areas. Electrophysiological studies (recordings of cells' response to various visual stimuli) and lesion studies (correlating brain damage with behavioral deficits) show that each one of these areas specializes in processing and representing a limited range of visual features. To give a few examples: in the early stages of the primary visual pathway (in the lateral geniculate nucleus of the thalamus) cells are organized in a two-dimensional map, and respond selectively to dots appearing in particular locations in the visual field. In the next visual area along this pathway (area V1) many neurons respond selectively to lines of a particular orientation in a particular spatial location. From there the information bifurcates into several distinct neural structures. Some of the information goes to areas responsible for analyzing various parameters of spatial organization and stimulus-motion, such as speed and direction of movement (e.g., in the MT area). A separate stream of information continues to area V4 which carries information about the color of the stimulus (probably among other parameters). In another area, or rather a cluster of distinct areas (inferotemporal), neurons are found to carry information about highly complex objects, such as faces and hands.

Current knowledge is far from complete, and there are still controversies about fundamental issues. Nevertheless, it is clear that each visual area carries information only about a narrow range of parameters (although there is some overlap and redundancy). For example, cells in area MT respond to motion, but display no significant responsiveness to the color of the stimulus, or to its complex shape. Cells in the inferotemporal areas may respond selectively to faces or hands, but not to the spatial location or motion of the stimulus, nor to simple features such as lines or dots. And cells in V1 that respond to line-orientation are insensitive to complex shapes or to dots.

Despite the specialization of distinct visual structures in different types of visual feature, there is no known neural structure which integrates all this information together. No single area was found to comprise neurons that respond selectively to motion, color, dots, line-orientation, and complex shapes. This implies that phenomenal events are not subserved by any particular structure, but must be spread over many different neural areas.

Furthermore, the visual field often comprises features that are processed and represented outside the visual system. Visual phenomenal features often contain lexical and semantic information (a written set of symbols may appear as a particular word), emotional information (an object may appear as frightening, beautiful, disgusting), various cognitive meanings, a spatial location, and associations with memories. All these are known to be processed and represented in various structures, mostly non-visual, throughout the brain. Moreover, phenomenal visual features are often integrated with phenomenal features of other modalities. A visually experienced drum may appear as the source of a drumming sound, and as the same object felt tactually. Non-visual sensory information is processed and represented in non-visual sensory structures.

In contrast with this distributed picture of neural representations, the phenomenal field constitutes a single unified scene. Phenomenal dots, lines, complex shapes, colors, motion, emotional value, cognitive meanings, and non-visual features such as sounds, are integrated in complex ways, often by merging together or "coloring" each other, within a single landscape. It seems, therefore, that phenomenal events are distributed over a large multiplicity of areas in the brain. This suggests that phenomenal properties express global patterns of activity across diverse neuronal areas. They apply globally to large areas in the brain (or to neural activity occurring in them).

Phenomenal Properties as Irreducibly Global

The conclusion that phenomenal properties are likely to apply globally to large parts of the brain sheds light on the possible role which they might play. Since the phenomenal field expresses the overall, integrated activity across many neural structures, it is a plausible hypothesis that it has the following two interrelated roles: first, to integrate the information distributed throughout these structures; and second, to coordinate between the activities of remote neurons or neural structures, by modulating neural activity through top-down causal signals that are sensitive to this integrated information. Due to their sensitivity to global information, these influences can impose a unified overall organization upon neural activity. It can be said, in short, that the phenomenal acts as a global overseer.

This conclusion is especially appealing in light of the complexity of the brain. The brain constitutes an astonishingly complex neural network, comprising more than 100 billion neurons, each one forming many thousands of connections, or synapses, with other neurons (see review in Kandel and Schwartz, 1985, chapters 1–12). Neurons typically form complex patterns of connectivity with hundreds or thousands of neurons of a variety of types. The behavior of each neuron is in itself extremely complex, much more than the simple "summarizer of inputs" portrayed in old introductory textbooks. It is governed by a large number of factors, such as the temporal pattern (e.g., frequency) of input-signals, interactions between different input-signals, the relative location on the neuron of each of the hundreds or thousands of input-synapses, more than ten distinct types of ionic current across the membrane, various chemicals secreted into the extra-neuronal fluids, various sorts of interaction with neighboring neurons, and a variety of long- and short-term effects of previous activations (ranging from milliseconds, to hours, to more or less permanence) [McCormick, 1990]. It is likely that many other factors are still unknown. These factors are not merely part of a stable silent background. They vary dramatically across neural structures and even across neighboring neurons, they change and interact constantly, and play crucial roles in actively modulating the neuron's behavior.

Despite this virtually unimaginable complexity, and despite the fact that even simple tasks involve many millions of neurons, the overall cognitive system displays consistent input-output relationships, a capacity to perform demanding tasks, adaptability to changing conditions, and an ability to learn and improve. This naturally raises the issue of how such a complex substratum can possibly give rise to the nicely-organized overall activity of the brain.

An attractive answer is that phenomenal properties play the role of integrating the activity in millions or billions of neurons. Due to their global nature, phenomenal properties carry information about the global pattern of neural activity, and so their effects on neurons can be used as top-down signals, in which information about the overall state of the system is used to modulate local neural activities.

This hypothesis would allow the neural microcircuitry to be only loosely organized. The task of neural microcircuits would be to give rise to patterns of activity that are only approximately right, while their precise fine-tuning would be performed by phenomenal properties. This is consistent with the apparent laxity that is commonly found in neural microcircuits (White, 1989). In the cortex, for example, neurons tend to "spray" their target-areas with output-synapses in an apparently indiscriminate manner. It appears that neurons project to populations of targets rather than to specific individual neurons, and that there are frequent spill-overs to neighboring targets. Admittedly, this laxity may be only apparent, and may merely reflect current ignorance of neuronal connectivity. But the globalist can in principle exploit this apparent gap in knowledge, and fill it with phenomenal properties.

It is important to see that in order for phenomenal properties to carry out this integration task, their causal powers must not be equivalent to the causal powers of their neuronal substratum. Otherwise, their contribution to the function of the brain would not differ from the contribution of the neural activity subserving them. Hence, phenomenal properties should not only apply globally, but their causal powers should not supervene on the causal powers of less global properties. They must be, in other words, *irreducibly global* properties.

Their global nature would make phenomenal properties especially suitable for the role of a global overseer. It can be postulated that specific types of complex patterns of neuronal activity give rise to instances of specific types of global properties. Consequently, the mere fact that a given global property is instantiated would already reflect the existence of a specific type of neuronal pattern of activity. This would amount to the integration of information from millions of neurons without having to go through low-level computations. The hypothesis is nicely consistent with the finding that, in some neural structures, sensory information is expressed in global patterns of activity across many millions of neurons (Freeman, 1991). Correspondingly, a global property could exert a multiplicity or correlated local influences on the many neurons subserving it. Such an effect would reflect the overall state of the system, and could therefore be used to impose overall patterns of organization on neural activity. In contrast, neural systems without global properties are likely to require much more complex integration mechanisms, since they need to build the overall picture from local bits of information distributed across vast numbers of neurons throughout the brain. Likewise, the regulation of large patterns of activity needs to be composed of a large number of local computations.

This is not to say that localist models of the brain are impossible. In fact, many ingenious models have been developed in the last decade, particularly within the connectionist approach, using statistical and other properties of neural networks in order to perform integration tasks (Rumelhart and McClelland, 1987). So far models have been commonly based on biologically unrealistic assumptions and over-simplifications, and limited to performing relatively simple tasks. But even if realistic localist mechanisms are possible, global properties—if they exist—are probably preferable from an evolutionary point of view. By bypassing the need for complex integration mechanisms, irreducibly global factors are likely to allow for simpler and much more lax neural microcircuits, and thus ones which can be developed more easily through evolution. Of course, local processes might still be an important part of the story, but not the entire story. Indeed, *non*-conscious neural activity might still be purely localist in nature (although, alternatively, it is possible to regard a non-conscious neural activity as an isolated island of consciousness that is not integrated with the rest of the person's consciousness).

Neural-Phenomenal Interaction

If indeed phenomenal properties are global properties which exert a top-down modulatory effect on their neural substratum, the question arises as to how exactly they are supposed to interact with this substratum. To answer this question the globalist would have to specify the law-like relationships between neuronal conditions and global properties. It is likely that the neuronal conditions which give rise to types of global properties are patterns of electric or chemical activity. Since the purpose of the discussion is only to outline some reasonable constraints on a globalist model, it would be inappropriate to be committed to any specific speculation. But it is worthwhile to note that neurons are known to display a variety of temporal, spatial, and causal patterns of activity, and there are ample candidates for the globalist to choose from. One possibility, which has been suggested in the literature (Kulli and Koch, 1991), is that phenomenal properties depend on synchronized patterns of neuronal firing. This is consistent with recent findings that, in several areas of the brain, different neurons responding to related stimuli fire in synchronization with each other (see review in Mitchison and Miall, 1990). It is a plausible speculation that these synchronized patterns determine, nomologically speaking, the instantiation of phenomenal properties, as well as the way they are merged together to form complex phenomenal features.

As for the opposite direction of the interaction, of the phenomenal on the neural, the globalist would have to specify the phenomenal conditions which, as a matter of law-like relations, produce given types of neuronal effects. This would presumably involve specifying a neural mechanism which is the recipient of the globalist effect, that is, whose states vary in correlation with phenomenal conditions. Since, as we will see later, global-neuronal effects may apply to microscopic quantum systems, this mechanism is likely to be on the order of magnitude of no more than a few molecules.

In order for such tiny mechanisms to substantially modulate populations of neurons—which are systems of several orders of magnitude larger—the globalist might stipulate that the desired global-neuronal effect is a simple sum of a myriad of tiny effects that add up to a significant size by the sheer force of their number. It seems more plausible, however, that the substantial size of the global-neural effect is a result of an amplification of a relatively small number of local effects. A smaller number of effects might require a less complex and more controllable mechanism, and consequently one that is more easily producible through evolution.

One attractive candidate for the locus of the global-neuronal influence is one or more of the several types of ion-channel that are found in the neuron's membrane (see reviews in Kandel and Schwartz, 1985, chapters 8–10; McCormick, 1990). An ion-channel is basically a pore in the external membrane of the neuron, which can be either open or closed. In its open state it allows the flow of ions from the extra-neuronal fluid into the neuron, or vice versa. Because of the difference in ion concentration and in electric charges between the two sides of the membrane, the opening of ion-channels normally results in a flow of ions into or outside the neuron. Ion-channels are typically selective in allowing through only ions of specific types. Although much of the function and structure of the different ion-channels is not yet clear, it is known that they have very significant short-term and long-term effects on the neuron's behavior: on its firing-threshold, on its mode of firing (e.g., single spikes versus bursts), on its firing rate, on the onset of firing, on adaptation to repeating input signals, and on the discharge of neurotransmitter in the synapse. In fact, ion-channels underlie the propagation of signals along the neuron.

An ion-channel is a structure made of several macro-molecules, which can be triggered in various ways to assume alternative spatial configurations. Microscopic fluctuations in local conditions may induce these molecules to pass from their open to their closed states, or vice versa. If global properties can exert tiny influences on the channel molecules, or on the relevant local conditions around them, they could influence the number of open channels of a given type. Depending on the type of channel, they could thus modulate neural activity in various ways. They could lower or raise the neurons' sensitivity to input signals, and thus their general activity level. They could modulate the neuron's firing mode and rate; an influence which the globalist could relate to the possible speculation mentioned above that global properties are nomologically related to patterns of synchronous firing. They could also affect calcium ion-channels in the synapse. The concentration of calcium ions in the synaptic bouton is a major factor in controlling the release of neurotransmitters (the chemical substances which mediate the passage of signals from one neuron to another). By modulating the number of open calcium channels, global properties could control the amount of neurotransmitter release, and consequently the strength of output signals produced by the neuron.

The magnitude of the global-neuronal effect should depend on the number of channels which the global properties are able to influence at any particular moment. But it may be postulated that even relatively small effects are amplified through various mechanisms. For example, a small contribution to the activity level of a neuron, if delivered at the right moment, may push a neuron over or under its firing threshold, and thus amplify a tiny contribution into the much larger effect of firing. A similar threshold-effect might be stipulated to occur in the synapse, in which neorotransmitter molecules are normally released in a small number of packets containing several thousands of molecules each. If at the right moment global properties cause tiny changes in the calcium ion concentration (e.g., by opening calcium channels), they might push the neurotransmitter release mechanism over its threshold, make it release an additional packet, and thus add a substantial amount of strength to the output signal. Furthermore, given our vast ignorance of the exact pattern of connectivity between neurons, it can be stipulated that small populations of neurons form mutual connections which are capable of amplifying small increases in activity through mutual excitation. This would be consistent with the prevalent view, according to which many parts of the brain, and in particular the cortex, are organized in modules of heavily interconnected populations of several hundreds or thousands of neurons (Eccles, 1984).

Whatever the exact locus of the global effect, in order for it to be capable of imposing proper organization on neural activity, it must be assumed to apply selectively or differentially to neurons of different type, function, and location. Clearly, indiscriminate effects are not likely to be functional. Here again the field is open for speculation. Neurons exhibit a wide variety of morphological and chemical parameters, and it is plausible to assume that global properties act differently on different types of entity. For example, excitatory and inhibitory neurons (i.e., ones whose output signal tends to excite or to inhibit the neuron receiving the signal) use different types of neurotransmitter, and have different types of synapse (White, 1989, chapters 1–2). It can be assumed that global properties have different nomological relations to different types of synapse or neuron, and that these differences are exploited by the brain.

It goes without saying that future findings might make these specific mechanisms of global-neuronal interaction implausible. But it should be emphasized that they are offered here only as examples designed to show that globalism is not inconsistent with current neurobiological knowledge. Indeed, alternative globalist models may be just as plausible as the one here proposed. For example, global properties may modulate neural activity not through action on ion-channels, as suggested above, but rather through action on some of the many enzymes which play crucial roles in regulating neuronal functions. The multiplicity of alternative globalist models does not contradict our point, but on the contrary, it shows that from the point of view of modern neurobiology globalism is a plausible hypothesis which deserves to be investigated.

Conclusion

To sum up, the globalist view can be made consistent with our current neurobiological knowledge if it portrays phenomenal properties as irreducibly global properties associated with large patterns of neural activity across the brain. By being associated with patterns of neuronal activation they carry information about the overall state of the system. Their role may therefore be that of integrating information and imposing overall organization on neural activity. This role may be executed through small local fine-tuning influences on microscopic mechanisms that are known to regulate patterns and levels of neural activity, such as ion-channels. Such local effects may add up to substantial modulation, possibly through some amplification mechanisms.

Part 2: Globalism and Quantum Mechanics

In accord with our aim to avoid the epiphenomenalist conclusion that consciousness is simply the music emanating from a neuronal instrument, it was concluded in Part One that irreducibly global properties might exert a causal influence on micro-events in the brain, an influence which is not contradicted by our current best theories of brain activity. The question we must now examine is whether this is consistent with our current best physical theories. There are two basic issues here. First, it may be objected that the physical world consists only of entities and properties characterizable in a localist manner. Second, it may be objected that the domain of entities and properties which are characterized in this manner is causally closed, contrary to our interactionist position.

The Source of Localist Prejudices

We will begin by considering the first line of objection. Localism, as an ontological thesis, had its home in the world-view of classical physics. An integral part of the classical mechanical world-view is the claim that physical reality can be analyzed into distinct parts, each having a definite intrinsic nature, and that the properties of composite physical systems are reducible to the intrinsic properties of the parts. For example, in classical physics the gravitational attraction between two material bodies is a relational property which is a function of their respective masses, and the distance between them (to be specified in terms of their respective positions in absolute space). The view that (1) the world can be analyzed into parts, each with definite intrinsic properties, and (2) that the properties of composite systems are reducible to the properties of the components of such systems, will be termed *Micro-Macro Reducibility Thesis* (MMT). According to MMT, the transition from the microphysical domain to the macrophysical domain should be a conceptually continuous one—macroscopic entities and properties should be reducible to microscopic entities and properties. This implies that there are no irreducibly macroscopic, or global, properties.

For example, the crowning achievement of statistical mechanics in the 19th century was a reduction of such macroscopic properties as heat, entropy, pressure, etc. (found in phenomenological thermodynamics), to the dynamical properties of a myriad of microscopic particles, viewed as components of macro-systems exhibiting the thermodynamical properties. For another example of MMT, given the microphysical ontology of small impenetrable particles characteristic of classical physics, it might seem, on the face of it, that the macrophysical property such as hardness is irreducibly macroscopic—that is, it is a property so radically different from those properties such as position, velocity, mass, etc., that the microphysical theory simply could not account for it. This, however, was not the case, for hardness could be accounted for in terms of the strengths of molecular bonds (in a way too well known to be worth recounting).

Arguably, then, it is the classical mechanical world-view that is the source of localist prejudices, and hence ontological hostility that there exist irreducibly global properties of brains. However, classical physics is not our current best physical theory. A question that naturally arises is whether contemporary physics provides evidence for or against localism. In this connection, two points will be made. First, the quantum theory, from the standpoint of its usual or orthodox interpretation, seems to require violations of MMT. Secondly, there seems to be some empirical evidence against MMT, and this evidence appears to provide support for globalism. Moreover, this evidence appears to be relevant to the criticism that the domain of localist entities and properties is causally closed. It must be emphasized here that we do not intend to provide a specific quantum mechanical account of the nature and dynamics of causally efficacious, global properties of the brain. We merely wish to establish that the existence of such global properties is part of the ontological perspective of the quantum theory (one of our best current physical theories), and hence reference to such properties may be a legitimate part of the scientific enterprise.

Global Properties and Measurements in Quantum Mechanics

From the standpoint of the orthodox interpretation of quantum mechanics, the claim that there exist irreducibly global properties gains support from the standard conception of quantum states, in particular from an examination of the orthodox resolution of *the measurement problem*. In general, when measurements are performed, quantum systems undergo discontinuous, inherently stochastic non-unitary changes of physical state. Such state transitions are sometimes referred to under the rubric of *the collapse of the wave-packet*. It is well-known that these state transitions cannot be governed by the time-dependent Schrödinger Equation,[3] which is the master dynamical equation of quantum mechanics. Such state transitions are to be governed instead by a very different rule: the Projection Postulate.[4] Obviously, since

[3]The Schrödinger equation is the master dynamical equation of elementary non-relativistic quantum mechanics. In these terms the time-development of the state vector Ψ is determined by: $\mathbf{H}\Psi = ih\ \partial\psi / \partial t$, where $\mathbf{H}$ is the Hamiltonian operator and h is Planck's constant divided by 2π.

[4]According to the Projection Postulate, if a *measurement* of observable A on a system Q yields eigenvalue a_i, then the state of Q *immediately after measurement* is the eigenstate Ψ_i of A corresponding to eigenvalue a_i (ignoring degeneracy). The measurement result probabilities are

measurement interactions are a proper subset of the set of physical interactions (all other interactions being in the domain of the Schrödinger Equation), some means are required to differentiate measurement interactions from other physical interactions. In the history of orthodox quantum mechanics, two main strategies have emerged to accomplish this end.

First, there is the *mentalist strategy*. Here the claim is made that measurements take place whenever a conscious observer interacts with a quantum system. The reduction of the wave packet, according to this theory, results from the intervention of consciousness in the physical world—so when the experimenter looks at Schrödinger's cat, for example, it is projected out of limbo and into either the "live cat" or the "dead cat" state. On the face of it, it looks as though this option goes against the spirit of physicalism. Some commentators have suggested that this appeal to consciousness in the context of the quantum measurement involves the adoption of some form of Cartesian dualism (for example, Earman, 1986, p. 223). In the context of the present globalist proposal, such claims may have to be reassessed. Although historically it has been assumed that a measuring consciousness must be some non-physical *entity*, this need not be so. What is required is that it is a phenomenon that is not itself a quantum system, but which is able to interact with such systems. As standardly interpreted, the effects of such interactions are inherently stochastic, with the probabilities that various results will be found upon measurement being determined by the Born Rule. Such a phenomenon could nevertheless be a global *property* of the brain, rather than a non-physical entity.

The second strategy, clearly physicalist in spirit, involves an explicit abandonment of MMT. It construes a measurement in terms of interaction not with consciousness, but rather with macroscopic, classically describable objects. A version of this strategy can be found at various points in Bohr's writings, but a particularly eloquent statement can be found in the work of David Bohm. Bohm observes that macroscopic objects are classically describable. Measurements (and hence reductions of the wave packet) take place when macro-systems (classically describable) interact with micro-systems. As Bohm (1989, p. 585) puts it:

> We may give as an example the usual practice in science, whereby one obtains data from meter readings, spots on a photographic plate, clicks of a Geiger counter, etc. All these objects and phenomena have the property of being classically describable. A little reflection will convince the reader that all observations ever made in science have employed at least one such classically describable stage.

---

given by the Born Rule which states that if the quantum state of a system Q is: $\Psi = \Sigma_i c_i \phi_i$, where $\{\phi_i\}$ is a complete orthonormal set of eigenstates of observable A with corresponding eigenvalues $\{a_i\}$, then $|c_i|^2$ gives the probability that *upon measurement* the eigenvalue a_i will be found.

On this view, then, there is a difference in kind between the interaction of a photon and an electron (micro–micro) and the interaction of a photon and a Geiger counter (micro–macro)—only the latter will be a measurement interaction! If this account is correct, then the property of being classically describable will be an irreducibly global physical property of macro-systems, contrary to the requirements of MMT and localism. It follows that both of those orthodox interpretations are consistent with the existence of irreducibly global properties that influence quantum systems: the property of being a conscious phenomenon, or the property of being classically describable.

A Second Case for Global Properties in Quantum Mechanics

While causally efficacious global properties play a role in the account of the quantum measuring process, it may still be argued that the case for them rests on claims about the orthodox interpretation of quantum mechanics which may turn out to be false or ill-conceived. As it happens, however, there is another empirical issue which points in the direction of global properties. The issue centers on a physical situation, described by J.S. Bell, which may be analyzed in quantum mechanical terms. The quantum theory makes definite predictions (verified by experiments) concerning what should occur in this situation, and these predictions are not compatible with localism (see details in Teller, 1989).

In the experimental situation envisaged by Bell, pairs of particles (of an appropriate kind) are prepared at a common source in a special type of physical (quantum) state.[5] These particles are then allowed to travel in opposite directions from the source. When they are separated by a macroscopic distance (usually several meters), each particle in the pair is subjected to a measurement to determine values for a special quantum mechanical property called *spin*. This property is always measured along some specific direction or axis—so we are concerned with measured values for "spin along direction **a**," "spin along direction **b**," and so on. In the system of units usually selected, the measured value for spin along any given direction is either +1 or –1.

Consider a given pair of such particles, labelled X and Y. On the basis of the quantum state in which the particle-pair has been prepared, quantum theory predicts that if spin along **a** is measured on particle-X and the value found is +1, then if spin along this same direction is measured on particle-Y, the value –1 will be found. Thus, a measurement of spin along some direction on one particle can be used to determine the value of spin along that

[5]This state is called the singlet spin state $\Psi = 1/\sqrt{2}[\varnothing_1(+1)\varnothing_2(-1) - \varnothing_1(-1)\varnothing_2(+1)]$. For a relatively accessible presentation of the mathematics of the Bell argument, as well as a discussion of the experimental situation, see W.D. Sharp and N. Shanks, 1985.

same direction on the other member of the pair. This is known as the strict anti-correlation feature of the quantum singlet spin state. More generally, when experiments are performed on large numbers of similarly prepared pairs of particles to determine values for spin along **a** on one particle and spin along some other direction **b** on the other particle, statistical correlations among the measured values are observed. According to quantum theory these statistical correlations will be a function of the angle between the direction along which spin is measured on one particle and the direction along which spin is measured on the other. Hence, the probability for finding a specific result for a measurement of spin along **a** on one particle, in an appropriately prepared pair of particles, depends on the direction **b** along which spin is measured on the other. Experiments have been performed and these provide an impressive confirmation of the correlations predicted by quantum mechanics. (For details on these experiments, and for a discussion of the philosophical issues here, see Cushing and McMullin, 1989.)

It is very natural to believe that these quantum statistical measurement-result correlations result from localist processes that mediate between those two spacelike separated particles; for example, that some causal signal traveling to the two particles, or from one particle to the other, is responsible for correlating their measured spins. How else would one particle "know" about the spin of the other? In other words, it is tempting to think that these correlations must be understood in terms of the *properties of the parts* of the composite systems consisting of two spacelike separated particles and two spacelike separated spin-measuring devices. Such properties would have to be thought of as *hidden variables*—factors not mentioned in the quantum description of the world. These expectations are grounded in MMT.

It was J.S. Bell's crowning achievement, however, to show that the price to be paid for any such localist analysis of the correlations in question was a violation of the Lorentz invariance (no-action-at-a-distance) requirements of the special theory of relativity (another of our current best physical theories). What Bell did was to show that the localist[6] who respects the Lorentz invariance requirements of the special theory, is constrained in the analysis of the correlations by a mathematical inequality—*Bell's inequality*—and this inequality is subject to substantial violations by the predicted (and confirmed) quantum mechanical correlations. In this case the price of localism is scientifically unacceptable action-at-a-distance!

Doesn't the Bell argument show that quantum mechanics is committed to action-at-a-distance! Happily no! The quantum state relevant to the predic-

[6]In the present context, the theorist who believes that the particles—and possibly the measuring devices—have definite intrinsic properties which determine the individual measurement outcomes in Bell-type experiments, which in turn give rise to the observed correlation statistics.

tion of the Bell-type correlation statistics does not in itself imply that there exist localist properties which underlie the predicted correlations (statistical relations). Instead, the evidence seems to point most strongly in the direction of the view that the measurement-correlations are irreducibly global properties of quantum systems that do not supervene on localist properties. The quantum systems in question have the property of correlated-spins-probability, and this property cannot be broken down, so to speak, to localist factors. This is a conclusion that has occurred to a number of quantum theorists, and most notably, Teller (1989).

We conclude this section by noting that both from the standpoint of the quantum mechanical account of the measuring process and from specific predictions and experiments performed, a case can be made for the existence of irreducibly global properties. MMT and localism fail in the context of quantum mechanics.

Global-Local Interaction and Quantum Mechanics

So far we have dealt with the first possible criticism against the globalist view of phenomenal properties. We have established that, contrary to this objection, the existence of global properties is not inconsistent with the deliverances of our current best theory of the microcosm. The second issue concerns the effects which such hypothetical global properties might have on microscopic events within the brain, such as the opening or closing of ion-channels in the neuron. The question is if current physics is consistent with the idea that global properties might affect microscopic events in ways that are beneficial to brain activity. We have no concrete proposal for a precise mechanism, nor indeed do we have a general theory. Instead we shall confine our remarks to speculations about the *possibilities* for such causal effects. Our speculations will rely on the fact that our current best microphysical theories are indeterministic.

In the history of philosophy there is a long tradition in which attempts have been made to exploit indeterministic physics as a means to the explanation of how consciousness may bring about physical effects. It will not go amiss to briefly mention one such proposal.

The physicist A.H. Compton attempted to exploit quantum mechanical indeterminism to model human decision making. Popper (1972, p. 227) comments on the proposed device as follows:

> It consists of an amplifier which amplifies the effect of a single quantum jump in such a way that it may either cause an explosion or destroy the relay necessary for bringing the explosion about. In this way one single quantum jump may be equivalent to a major decision.

As Popper points out, this is not a very convincing model of rational decision-making as opposed to that species of decision-making where coins are tossed. The model is primarily defective in that it attempts to base the decision-making process directly on random quantum mechanical events. In this way the decisions made will be as stochastic as the micro-events which cause them.

It is our contention that there is a possibility here which has not been exploited. If, as has been argued in Part One of this paper, phenomenal properties (consciousness) may be irreducibly global properties of the brain, then they could be seen as performing measurement interactions with the quantum systems underlying their neural substratum, and hence as influencing brain states. However, in view of the comments made in connection with Compton's proposal, the causal effect of such a global property should not be entirely random—for if it is, it is hard to see how the appeal to any such property could play much of a helpful role in the analysis of brain function. In particular, it is hard to see how it could play a role in the integration of information and imposing overall organization, as proposed in Part One of this paper.

However, it is possible that consciousness is capable of influencing micro-events in the brain not through random processes, but rather by *manipulating* the occurrence of micro-events at the quantum level. Consciousness, qua global property, may not only perform measurement interactions, but also determine the results of measurements—and hence the physical states—of the microscopic components of the brain. In other words, consciousness would not only be responsible for the reduction of the wave-packet, as in orthodox quantum mechanics, but it would also be able to influence the result of the reduction. The statistical features of the quantum mechanics of the brain would then arise because the way in which consciousness does this is unknown in individual instances.

There are clear constraints on such an account of the role of consciousness. Such manipulation would have to be consistent with the probability distributions predicted by quantum mechanics for events happening to the relevant micro-systems. Unless quantum mechanics is false, these constraints cannot be violated. It should therefore be stipulated that global properties of the brain—i.e., consciousness—may manipulate the occurrence of individual quantum-events in the brain within those probabilistic constraints. The effects of such manipulation are then hypothesized to be amplified into events with consequences for the states of neuronal systems—such states being the subject matter of the above neurobiological discussion.

A simple analogy may explain how such manipulation may be consistent with a well-defined probability distribution. Consider coin tossing. For creatures like us, who have only an imperfect knowledge of the initial states of tossed coins, the results appear random, with a probability of 0.5 that on any given toss the outcome will be a head. From the standpoint of classical physics, however, if the initial state of the coin was known with perfect precision, then the precise outcome could be predicted with certainty. It is not hard to imagine a Casino of the future which has a deterministic coin-tossing machine such that the observed frequency of "heads" is 0.5 in a long series of trials, but where the outcome on any *particular* toss is under the control of the management. Thus, the management may manipulate some individual events in a favorable way (e.g., favor one gambler on the expense of another), if it is careful to comply with the general constraints of the probability of coin-tossing.

There is a more serious issue, however, concerning the very possibility of consciousness manipulating seemingly stochastic quantum events. What about the "no hidden variables arguments?" (For details on hidden variables see Jammer, 1974, chapter 7). Do these rule out such manipulation (or determination) of the results of individual measurements—and hence physical states—of the quantum components of the brain? The answer is not clear-cut. It is true that there exist quantum states—such as the singlet spin state discussed by Bell—for which hidden variables seem to be ruled out. So there appear to be certain quantum systems, some of whose states do not admit of a hidden variables analysis. However, as Bell himself points out, there exist quantum systems for which it is a trivial matter to produce hidden variables theories. In his classic discussion of the EPR paradox, he actually provides a hidden variables account of spin measurements on a single particle. He also provides a local hidden variables account for a restricted class of spin correlations for systems in the singlet spin state (the case where spin is measured along the same axis on each particle). What Bell's argument shows you cannot have is a hidden variables theory for *all* of the statistical predictions made relative to *every* quantum state found in quantum mechanics. The no hidden variables arguments do not rule out the less ambitious program of providing hidden variables accounts for restricted types of measurement relative to certain quantum states (a set of states which will be a proper subset of the set of quantum mechanical states).

To this end, it is worth noting that investigators of stochastic electrodynamics have succeeded in reproducing the correct quantum statistics for an impressive number of quantum systems (blackbody radiation, stability of the ground state of the hydrogen atom, the harmonic oscillator, van der Waals forces, etc.) using Newton's equations for particles and Maxwell's equations for the electromagnetic field. The ontology is both determinate (simultaneous exact values for all quantities at all times) and deterministic. So there actually exists a restricted class of quantum systems which admit of a consistent hidden variables analysis (Boyer, 1975, 1981). The general "no hidden variables" theorems do not, therefore, rule out the possibility that consciousness is capable of manipulating the results of individual measurements for certain classes of quantum systems. The sort of determination of individual measurement results—and hence physical states—of certain quantum components of the brain, of the kind which we are proposing is not, therefore, obviously inconsistent with the statistical predictions of our current best microsphysical theory.

Conclusion

We conclude that a "property-dualist" theory which views phenomenal properties as irreducibly global properties need not contradict either current knowledge of the brain or contemporary physical theories. The view that consciousness is to be understood in terms of global properties which apply to brains is consistent with contemporary science. These properties might be responsible for the reduction of the wave-packet, and also influence the results of the reduction, in many of the myriad of quantum systems out of which the brain is constituted. Thus, they may exert microscopic influences on various physical parameters of micro-systems, presumably in accordance with some still unknown law-like relation. This influence is constrained by the probability distributions predicted by quantum mechanics, and is limited to those micro-systems in which hidden variables are permissible. Since global properties are associated with (and apply to) overall states of brain structures, their influence carries information about the overall state of the system. This information can be exploited by the brain, in an evolutionarily advantageous way, in order to integrate information across neural structures and to coordinate between neural activities. To do this, the brain probably amplifies the microscopic global influence through special amplification mechanisms (such as ion-channels) which have developed through evolution, at least in part, precisely for this purpose.

All this leaves open the issue of whether there is a good reason—perhaps traditional anti-"physicalist" arguments—to posit global properties. It also leaves open the issue of the characteristics that need to be ascribed to phenomenal properties in order to account for their subjective and qualitative nature. And, of course, much more has to be added to the picture in order to make it a working model. But despite these gaps, to be filled in in the future, the point of the discussion is that "property-dualism" is not obviously unscientific, as it is very often accused of being. It remains to be seen whether this "dualist" possibility can form an avenue along which justified research programs can be developed. That is an empirical matter which cannot be settled by apriori considerations or a straightforward appeal to what is actually mandated by current neurobiology or physics.

References

Bohm, D. (1989). *Quantum theory*. Dover: New York.

Boyer, T.H. (1975). Random electrodynamics: The theory of classical electrodynamics with classical zero point radiation. *Physical Review D, 11*, 790–808.

Boyer, T.H. (1981). A brief survey of stochastic electrodynamics. In A.O. Barut (Ed.), *Foundations of radiation theory and quantum electrodynamics* (pp. 49–63). New York: Wiley.

Cushing, J.T., and McMullin, E. (Eds.). (1989). *Philosophical consequences of quantum theory*. Notre Dame: University of Notre Dame Press.

DeYoe, E.A., and Van Essen, D.C. (1988). Concurrent processing in monkey visual cortex. *Trends in Neurosciences, 11*, 219–226.

Earman, J. (1986). *A primer on determinism*. Dordrecht: D. Reidel.

Eccles, J.C. (1984). The cerebral neocortex. A theory of its operation. In E.G. Jones and A. Peters (Eds.), *Cerebral cortex* (volume 2, pp. 1–26). New York: Plenum Press.

Freeman, W.J. (1991). The physiology of perception. *Scientific American, 264* (February), 78–85.

Jackson, F. (1982). Phenomenal qualia. *Philosophical Quarterly, 32*, 127–136.

Jammer, M. (1974). *The philosophy of quantum mechanics*. New York: Wiley.

Kandel, E.R., and Schwartz, J.H. (1985). *Principles of neural science*. New York: Elsevier.

Kim, J. (1982). Psychophysical supervenience. *Philosophical Studies, 41*, 51–70.

Kulli, J., and Koch, C. (1991). Does anesthesia cause loss of consciousness? *Trends in Neurosciences, 14*, 6–10.

Lahav, R. (1990, February). *Neural correlates of visual consciousness*. Paper presented at Consciousness Within Science, First International Conference, University of San Francisco.

McCormick, D.A. (1990). Membrane properties and neurotransmitter actions. In G.M. Shepherd (Ed.), *The synaptic organization of the brain* [third edition] (pp. 32–66). New York: Oxford University Press.

Mitchison, G., and Miall, C. (1990). The enigma of the cortical code. *Trends in Neurosciences, 13*, 41–43.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review, 83*, 435–450.

Popper, K. (1972). *Objective knowledge*. Oxford: Clarendon Press.

Rumelhart, D.E., and McClelland, J.L. (1987). *Parallel distributed processing*. Cambridge: MIT Press.

Sharp, W.D., and Shanks, N. (1985). Fine's prism models for quantum correlation statistics. *Philosophy of Science, 52*, 538–564.

Teller, P. (1989). Relativity, relational holism, and the Bell inequalities. In J.T. Cushing and E. McMullin (Eds.), *Philosophical consequences of quantum theory* (pp. 208–223). Notre Dame: University of Notre Dame Press.

White, E.L. (1989). *Cortical circuits*. Boston: Birkhäuser.

Zeki, S., and Shipp, S. (1988). The functional logic of cortical connections. *Nature, 335*, 311–317.