Structural realism and the mind

Gordon McCabe

May 25, 2006

Abstract

This paper considers whether, and how, the mind can be incorporated into structural realism. Section 1 begins with some definitions, and briefly reviews the main problems which beset structural realism. The existence of the mind is proffered as an additional problem, to which the rest of the paper is devoted. Three different philosophies of the mind are analysed, beginning with eliminative materialism, which is briefly reviewed in Section 2. The identity theory of the mind-brain relationship is critically analysed in Section 3, and the notions of supervenience and emergentism are defined. In Section 4, the functionalist approach to the mind-brain relationship is introduced, and two specific functionalist approaches—the representational theory of the mind, and connectionism—are defined and appraised. It is argued that these approaches enable structural realism to be extended to include the mind. It is also argued that structural realism can be applied to the unconscious mind, and the paper concludes with the proposal that the distinction between epistemic structural realism and ontic structural realism is also valid in the case of the mind.

1 Introduction

The purpose of this paper is to discuss whether the existence of the mind can be incorporated into structural realism, and if so, how.

Let us begin, then, by reviewing some relevant definitions. A mathematical structure consists of a set, or a collection of sets, (a 'base' set or sets), equipped with properties, relationships, operations and distinguished elements. These entities are collectively required to satisfy a set of conditions, called the axioms of the structure. A base set equipped with such a structure is said to be a structured set. Whilst *structural empiricism* accepts that mathematical structures are used in physics to organize, explain and predict empirical phenomena, it holds that those structures do not characterize anything beyond the empirical phenomena. In contrast, *structural realism* asserts that, in mathematical physics at least, the domain of a true theory is an instance of a mathematical structure which lies beyond the empirical phenomena. It follows that if the domain of a true theory extends to the entire universe, then the entire universe is an instance of a mathematical structure. Equivalently, it can be asserted that

the universe is isomorphic to a mathematical structure (Tegmark 1998). One can refer to this proposal as 'universal structural realism'.[1]

Many authors take structural realism to assert that the physical world possesses a second-order structure, without specifying the physical sets and physical first-order properties and relationships which possess this structure (see Votsis 2003). This view is inspired by Bertrand Russell, who argued in *The Analysis of Matter* in 1927 that, at most, what we can *know* is the second-order structure of the world which lies beyond empirical phenomena, not the first-order physical properties and relationships which possess that structure. This position was dubbed Epistemic Structural Realism (ESR) by Ladyman (1998), and distinguished from Ontic Structural Realism (OSR), the assertion that the second-order structure of the physical world is all that exists.

There is, however, no need to accept this second-order approach to structural realism. Instead, one can assert that the physical world is a specific set, whose elements are equipped with specific, knowable first-order physical properties and relationships, which exist beyond the empirical first-order properties and relationships, and which satisfy the conditions of a certain second-order structure. Structural realism can be taken to assert that the physical world is a specific structured set, not merely a second-order structure.

Nevertheless, structural realism faces probably three main problems. Firstly, there is the so-called Newman problem. As Pooley puts it, "it is a theorem of set theory and second-order logic that any consistent proposition to the effect that a certain set of properties and relations exist, no matter what structural constraints are placed upon this set, will be true of any domain, provided that the domain has the right cardinality," (Pooley 2005). This suggests, then, that structural realism is trivial as an account of the reality beyond the phenomena; if a physical domain is a set of a specific cardinality, then what lies beyond the observable properties and relations of that domain will possess any structure of that cardinality. The empirical properties and relationships, and the structural conditions they satisfy, might indeed pick out a specific structure, but this is merely the structure of the empirical phenomena. Structural realism, then, faces the danger of either being trivial, or collapsing into structural empiricism. For structural realism to work, it must be possible to infer and define physical properties and relationships beyond the empirical properties and relationships, and for the conditions which these properties and relationships satisfy to pick out a specific structure.

Secondly, a theory with a fixed structure can possess more than one realist interpretation. For example, consider the many-worlds interpretation (MWI) and the Bohm-deBroglie interpretation of quantum mechanics. If there is more than one realist interpretation for a fixed theoretical structure, it arguably indicates that more than structure alone exists beyond the empirical data. This, of

[1]Whilst the definition of structural realism is most frequently expressed in terms of the set-theoretical, Bourbaki notion of a species of mathematical structure, one can reformulate the definition in terms of other approaches to the foundations of mathematics, such as mathematical category theory. In the latter case, one would assert that our physical universe is an object in a mathematical category.

course, is precisely the thesis of ESR, and Ladyman proposes OSR as a means of abstaining from further interpretation of a fixed structure. Personally, however, I would argue that whenever different realist interpretations of a physical theory exist, they each tacitly require either a modification or an addition to the existing structure of the theory. For example, DeWitt's version of the MWI tacitly invokes the notion of branching space-time, and this corresponds to a different mathematical structure for space-time than that assumed in conventional quantum mechanics.

Thirdly, there is the problem of so-called 'Jones under-determination'. This occurs when a theory has more than one formulation, each of which is empirically equivalent. (This should not to be confused with the 'under-determination of theory by data', which occurs when there are distinct theories, each of which is empirically equivalent). The problem for structural realism occurs when the different formulations of a theory involve different mathematical structures. For example, in the traditional formulation of canonical classical general relativity, the configuration space is the set of Rimennian metric tensor fields on a 3-manifold Σ, but in Ashtekar's 'new variables' formulation, the configuration space is the space of connections upon an $SU(2)$-principal fibre bundle over Σ. These are non-isomorphic structures. Whilst Schrodinger's wave mechanics formulation of quantum mechanics did, indeed, transpire to be isomorphic to Heisenberg's matrix mechanics formulation, such an outcome is highly untypical. Clearly, if there are non-isomorphic formulations of a theory, then a structural realist cannot say which of these structures the theory is committed to the instantiation of. It is true that the different formulations of a theory must share some common structure in order to be empirically equivalent, but the structure they share in order to achieve this, is the structure of the empirical phenomena, hence, once again, structural realism runs the risk of collapsing into structural empiricism.

These are serious, but well-known problems for structural realism. In this paper I wish to draw attention to an additional problem which, to the best of my knowledge, has gone previously unacknowledged. Given that structural empiricism does not purport to extend beyond the realm of empirical phenomena, it is not under any obligation to include the mental within its remit. In contrast, universal structural realism purports to have an unrestricted range of validity. Because universal structural realism claims that the universe is an instance of a mathematical structure, it is necessary to consider whether, and how, the mind can be incorporated within it.

There are three distinct positions in the philosophy of the mind which I wish to consider. Broadly, they can be defined as follows:

- Eliminative materialism: The brain exists, but the mind does not.
- Identity theory: The mind and the brain both exist, but the mind can be reduced to the brain in the specific sense that the mind can be identified with the brain.

- Functionalism: The mind and the brain both exist, and the mind supervenes on the brain, but the mind cannot be identified with the brain.

Whilst eliminative materialism and the identity theory of the mind-brain relationship provide no reason to modify or extend structural realism, I will argue that functionalism opens up some interesting new possibilities. Before we proceed, however, it should be emphasized that the primary purpose of this paper is to illustrate the possibilities for extending structural realism, and, as such, the approach adopted to the philosophy of the mind is highly selective and non-systematic.

2 Eliminative materialism

Eliminative materialism (EM) accepts that the mind cannot be reduced to the brain, but claims that the psychological terms used to describe the mind are terms which have no referents. Unlike the identity theory discussed below, EM does not claim that psychology and neurophysiology provide two different levels of description of the same system; rather, EM claims that the language of psychology does not provide a description of anything at all. Eliminative materialism obviously poses no problems for universal structural realism: if the mind doesn't exist, then structural realism merely has to incorporate the structure of the brain into its ontology. The universe is an instance of a mathematical structure, and each brain is merely an instance of a substructure in the universal structure. We therefore proceed straight to the more interesting case of the identity theory.

3 The identity theory

The identity theory asserts that the mind and the brain both exist, but the mind can be reduced to the brain in the specific sense that the mind can be identified with the brain. It is clear that there are very strong correlations between the mind and the brain, and research in neurophysiology has enabled the understanding of these correlations to become increasingly sophisticated. The identity theory suggests from this that we can re-define what mental properties, states and processes really are in terms of brain properties, states and processes.

If mental properties can be identified with brain properties, if mental states can be identified with brain states, and if mental processes can be identified with brain processes, then the existence of minds poses no problems for universal structural realism. If minds can be reduced to brains in the sense that mental properties, states and processes can be *defined* in terms of brain properties, states and processes, then because universal structural realism is capable of incorporating the existence of brains, it is automatically capable of incorporating the existence of minds.

However, does the identity theory provide a plausible account of the mind? It is difficult to understand how the subjective nature of experience, which

arguably characterises the mind, can arise from an objective physical system such as the brain. In fact, Colin McGinn (2004, p158) lists five characteristics of mental states, which purportedly prevent them from being identified with brain states. Mental states are:

- Unobservable — in the sense that they are not perceptible by means of the senses.
- Asymmetrically accessible — in the sense that the owner of a mental state has a kind of immediate access to it that other people do not.
- Subjective — in the sense that the nature of a mental state is knowable only from a single 'point of view'.
- Non-spatial — in the sense that mental states do not take up a well-defined region of space.
- Subject-dependent — in the sense that mental states only exist for a subject of awareness.

In a similar vein, those who endorse the notion that the mind cannot be identified with the brain, have a favourite mantra, in which, following Nagel (1986), they recite that there is something which *it is like* to be the subject of experience. The force of this claim seems to reside in the very fact that it makes any sense. The very fact that the reader knows what it is like to be a subject of experience, is the only thing which can make this a meaningful claim. Whilst this slogan is a good weapon against those who seek to eliminate the mind, it only entails that the mind cannot be identified with the brain if it can also be shown that there is nothing which it is like to be a brain. If the mind exists, if there is something which it is like to be a mind, and if the mind *can* be identified with the brain, then it follows that there is something which it is like to be a brain.

Let us consider the identity theory in terms of inter-theoretical reduction within science and physics. One type of reduction relationship between a theory T_1, and another theory T_2, occurs when T_2 represents a larger domain of phenomena than T_1, and either the laws and theorems of T_1 can be derived as special cases from those of T_2, or the laws and theorems of T_1 are at least approximated by a subset of what can be derived from T_2 (Butterfield and Isham 1999, Section 2.2). Thus, for example, special relativity represents a larger domain of empirical phenomena than Galilean relativity, but there is a limit $v/c \to 0$ in which special relativity approximates Galilean relativity. However, if we think of the mind as being theoretically represented in psychology, and the brain as being theoretically represented in neurophysiology, then these two theories are not related by this type of reduction relationship. Instead, according to the identity theory, neurophysiology provides a description of purportedly the *same domain* as psychology, but with a higher level of resolution, a greater level of detail. Moreover, in the case of a highly complex system, such as the brain, it may not be practically possible to use the higher level of resolution of the

reducing theory to re-derive the laws of the reduced theory; the problem may be computationally intractable. Thus, a second type of reduction relationship occurs when T_2 represents the same domain as T_1, but with a greater degree of resolution, and either the laws and theorems of T_1 could, in principle, be derived from T_2, or the laws and theorems of T_1 are at least approximated by a subset of what could, in principle, be derived from T_2.

In general, a theory which provides a higher level of resolution often does so by postulating entities or structures of a smaller size than those in a lower-resolution theory. Hence, reductionistic explanations often explain macroscopic phenomena in terms of the states and processes of much smaller constituent particles and fields. The best example is the reduction of thermodynamics to statistical molecular mechanics.[2] As a consequence of increasing the level of resolution, there is a one-many correspondence between the states of the reduced theory and the states of the reducing theory. The theory with a lower level of resolution provides a coarse-grained description, while the theory with a higher level of resolution provides a fine-grained description. Each state in the coarse-grained theory corresponds to multiple states in the fine-grained theory, and it is only the higher level of resolution provided by the fine-grained theory which enables it to discriminate between those different states. One conventionally refers to the states of the lower-level, higher-resolution theory as 'microstates'. One can partition the states of the lower-level theory into equivalence classes, consisting of those states which are indistinguishable at the coarse-grained level of resolution provided by the higher-level theory. It is conventional to refer to these equivalence classes as 'macro-states'. The macro-states correspond in a one-to-one fashion with the states of the higher-level theory.

Arguably, it is a necessary condition for a theoretical reduction relationship that the higher-level properties are *definable* in terms of the lower-level properties. For example, temperature in thermodynamics is definable as mean molecular kinetic energy in statistical molecular mechanics. Following Nagel, these inter-theoretical definitions are called 'bridge laws'. In a reductive explanation, the domain of enquiry in re-conceptualized in terms of the ontology of the lower-level theory, and in particular, the bridge laws are commonly taken to reveal the true identity of the higher-level properties. Thus, the temperature of a gas *just is* the mean kinetic energy of its molecular constituents.

So, is the mind-brain relationship a reduction relationship of the second type defined above? Certainly, a mental state is a very coarse-grained, inexact level of description, whilst a brain state is a fine-grained and exact level of description. Hence each mental state corresponds to multiple brain states, and one can partition brain states into mentally indistinguishable equivalence classes. Let us call these equivalence classes 'brain macro-states'. Similarly, there are brain macro-processes, which evolve one brain macro-state into another, in correspondence with the evolution of one mental state into another. However, the relationship between psychology and neurophysiology fails to satisfy the cited

[2]Note, however, that there remains a debate in the philosophy of science over whether this reduction is partial or complete.

necessary condition for a reduction relationship, that of *defining* the higher level properties in terms of the lower-level properties. The feeling of pain and the perception of colour, for example, are psychological properties which cannot be defined in terms of neurophysiological properties. A psychological property is unobservable, asymmetrically accessible, subjective, non-spatial and subject-dependent, while a neurophysiological property is "observable by means of the sense organs; it is accessible to oneself and others in the same way; it is objective in that it can be grasped from any point of view. . . ; it is spatially defined; it could, in principle, exist without being experienced by a subject," (McGinn 2004, p158).

In particular, the suggestion that mental properties be defined in terms of brain properties, should be rejected because the brain is spatial, but the mind is not; the two levels of description are non-comparable in this respect. The brain and the mind do not provide different levels of *spatial* resolution of a fixed domain. Although the one-many mapping between mental states and brain states entails that the brain provides a lower-level, higher-resolution description, it does not provide a re-definition of the terms contained in the higher-level mental description.

A more acceptable concept than mind-brain identification is the weaker notion of supervenience, which holds that any change in the higher level properties, states and processes of a composite system, must correspond to a change in the lower level properties, states and processes. The idea is that there can be no difference in the higher-level state of a composite system without a difference in the lower-level state, otherwise there would be a one-many correspondence between the lower-level states and higher-level states. In these terms, the mind clearly supervenes upon the brain: each brain state determines a unique mental state, and each change in mental state requires a change of brain state. However, supervenience does not require the higher-level description to be definable in terms of the lower-level description. In particular, supervenience does not require the properties of the higher-level description to be definable in terms of the properties of the lower-level description. When supervenience is combined with this claim of irreducibility, the resulting credo is often termed 'emergentism'. Emergent states and properties are higher level states and properties which supervene upon lower level states and properties, but which are not definable in terms of those lower-level properties.

Note, however, that there is a distinction between ontological emergentism and epistemological emergentism. The former claims that 'genuinely novel' properties exist at a higher level, whilst the latter merely claims that the human intellect may lack the imaginative capacity or conceptual resources to define the higher level properties in terms of the lower level parts, properties and relationships. Rejection of the mind-brain identity theory does not entail acceptance of ontological emergentism.

One can also introduce a distinction between ontological mind-brain identity and epistemological mind-brain identity. The latter asserts that our conception of the mind can be identified with our conception of the brain, and it is this version of the identity thesis which I have rejected. The ontological version

of the mind-brain identity thesis accepts that our conception of the mind may not be identifiable with our conception of the brain, but holds that that which appears to us as the mind, and that which appears to us as the brain, exist as the same thing. Epistemological emergentism and the ontic version of the mind-brain identity thesis are consistent with each other.

4 Functionalism

Functionalism accepts that the mind and the brain exist; it accepts that the mind cannot be identified with the brain; and it accepts that the mind supervenes on the brain. Functionalism claims that the mind is a set of functionalities and capabilities, at a higher level of description than the brain. Thus, although functionalism accepts that the mind cannot be identified with the brain, it still contends that the mind can be objectively characterised.

The notion that the identity of the mind is defined by a set of functionalities and capabilities, leads to the notion of substrate-independence, the claim that the mind could supervene upon multiple substrates, of which the brain just happens to be one example. Neurophysiology demonstrates how brain structure supports the functionalities and capabilities of the mind, hence neurophysiology demonstrates how the structure of the brain supports these mental structures. However, functionalism argues that there are multiple substrates which could support such mental structures, hence one cannot identify the mind with the brain.

From the perspective of structural realism, functionalism holds that the mind possesses a structure which cannot be identified with the structure of the brain. In other words, functionalism holds that the structure of the mind is non-isomorphic to the structure of the brain.

Functionalism holds that a mental state has functional relationships to perceptual stimuli, behavioural responses, and other mental states. i.e., a mental state maps current perceptual stimuli to behavioural responses and the next mental state. Thus, in mathematical terms, one can treat a mental state as a function

$$I \to S \times O \ ,$$

where I is the set of input states (the perceptual states), S is the set of mental states, and O is the set of output states, (the behavioural responses). By implication, the set of mental states is then a set of such functions, and this set presumably possesses some structure.

Functionalism, however, should not be conflated with 'behaviourism', which claims that mental states have no 'internal' content, and are *nothing but* maps between perceptual stimuli and behavioural responses. Functionalism claims that mental states are a lot more than such maps. We shall now consider two specific functionalist approaches to the mind: the 'top-down' approach provided by the conjunction of the representational theory of the mind (RTM) and the

computational theory of the mind (CTM), and the 'bottom-up' approach of *connectionism*. In particular, we shall attempt to highlight the structural aspects of these approaches.

The RTM attempts to provide a functionalist account of 'intentional' mental states. These are states, such as beliefs and desires, in which the attention of the mind is directed towards something, called the 'content' of the intentional state. The RTM claims that an intentional mental state is a type of functional state that involves a relationship between the thinker and the symbolic representation of something. If a thinker holds the belief that 'the cat is on the mat', then the thinker is held to be in one type of functional relationship to a symbolic representation of 'the cat is on the mat'. If a thinker holds the desire that 'the cat is on the mat', then the thinker is held to be in a different type of functional relationship to the same symbolic representation. Beliefs and desires differ by virtue of the fact that they cause different succeeding mental states and behavioural responses. The RTM considers mental processes such as thinking, reasoning and imagining to be sequences of intentional mental states.

At this point it is judicious to introduce the distinction between the conscious and the unconscious mind. There are, of course, various approaches to defining consciousness. For example, in the so-called 'higher-order' theory, a conscious mental state is a state which is accompanied by a higher-order mental state, an intentional state whose content is the first mental state. An unconscious mental state is then defined to be a mental state which is not accompanied by such a higher-order state (van Gulick 2004, Section 9.1). Our introspective knowledge of the mind is limited to knowledge of the conscious mind, and it is something of a cliche to say that the conscious mind constitutes only the surface of the mind. Perhaps, however, a more powerful analogy is to suggest that the conscious mind provides an interface to the unconscious mind, in the way that a web browser can provide an interface to a remote database. Processes such as the unconscious execution of familiar tasks, and the unconscious retrieval of facts from memory, suggest that there is something of a client-server relationship between the conscious and unconscious mind. The RTM may or may not be applicable to the unconscious as well as the conscious mind, but, at the very least, one can expect the unconscious mind to possess a structure, in the same sense that a relational database possesses a relational structure[3]. The unconscious mind must possess some sort of structure in order to perform the tasks which it does perform. We shall return to the unconscious mind later below.

Many advocates of the RTM claim that the mental representations which provide the content of beliefs, desires, and other intentional states, possess an internal structure. They hold that this internal system of representation has a set of symbols, a syntax, and a semantics, collectively termed the language of thought. There are rules for composing the symbols into expressions, propositions, and mental images, hence the content of an intentional state can be said to possess a symbol structure, and one might call this the infrastructure

[3]In computer science, the data in a relational database is a collection of 'tables', each of which comprises a set of n-tuples. Such n-tuples correspond, with some differences, to the extension of an n-ary relation.

of an intentional state. The computational theory of mind (CTM) is then the conjunction of the RTM with the claim that mental reasoning is the formal, syntactical manipulation of such symbols (Horst 2005).

In contrast with the CTM, connectionism does not attempt to model the mind by ascribing a structure to intentional states at the outset. Instead, connectionism adopts a 'bottom-up' approach, modelling the mind with 'neural networks', abstractions from the network of nerve cells and synapses in the human brain. In the connectionist approach, the foremost structures are those possessed by the neural networks, not those possessed by sets of intentional states, or those possessed by the symbols which purportedly compose the contents of intentional states.

A neural network consists of a set of nodes, and a set of connections between the nodes. The nodes in a neural network possess activation levels, the connections between nodes possess weights, and the nodes have numerical rules for calculating their next activation level from (i) the previous activation level, and (ii) the weighted inputs from other nodes. The pattern of connections and activation levels can be thought to provide a lower-level of description than that provided by sets of intentional states and symbol structures. However, under certain circumstances, the pattern of connections and activation levels is deemed to provide a representation of things, hence the connectionist models of the mind can still be subsumed under the aegis of the RTM. It should also be noted that connectionism (arguably) contrasts with the CTM in the sense that thinking can take place at a sub-symbolic level.

To answer the question posed at the outset, these functionalist accounts of the mind entail that the mind *can* be incorporated into structural realism. These accounts claim that the mind possesses a structure, albeit a higher-level structure than that of the brain. All functionalist accounts accept that the mind supervenes on the brain, but they reject the idea that the higher-level structure can be defined in terms of the lower-level structure. One can accept that the mind is essentially subjective, but one can still hold that it possesses a structure, a distinct structure from the structure of the brain. The structure may well be a non-spatial structure, which cannot be reduced to any of the structures which characterise the brain, but it is, nevertheless, the structure of subjective experience. Intentional mental states are not observable by means of the sense organs, are directly accessible only to their owners, and do not occupy space, yet they nevertheless possess a structure which functionalism and the RTM sets out to capture.

However, functionalism and the RTM do not capture all there is to capture about the mind. Whilst one can represent the second-order structure of the mind, just as one represents the second-order structure of the extra-mental world, such structure does not enable one to know what it is like to feel pain or see colour. Whilst the second-order structure of the mind can be known, we cannot identify the mind with its second-order structure. Just as structural realism accepts that there are known first-order empirical properties and relationships, the application of structural realism to the mind should accept that there are known first-order mental properties and relationships, such as the

qualia of mental states. The *quale* of a mental state is that state's feel, and one might claim that "it is the feels that make the mental states the mental states they are," (Lycan 1994, p322). The quale of a mental state can be treated as a first-order property of the state, and, in general, introspection provides us with knowledge of *conscious* mental states and their first-order properties.

To conclude, I propose that the distinction between epistemic structural realism (ESR) and ontic structural realism (OSR) is also valid in the case of the mind. The conventional application of structural realism asserts that there is a physical world which exists beyond the empirical phenomena, and this physical world possesses a second-order structure; analogously, the application of structural realism to the mind should assert that there is an unconscious mind which exists beyond the introspectively-known conscious mind, and this unconscious mind possesses a second-order structure. Having made this proposal, ESR would then assert that the second-order structure of the unconscious mind is all that can be known of it, and OSR would assert that the second-order structure is all that exists of the unconscious mind.

References

[1] Butterfield, J. and Isham, C.J. (1999). On the Emergence of Time in Quantum Gravity, in *The Arguments of Time*, J. Butterfield (ed.), pp111-168. Oxford: Oxford University Press.

[2] Horst, S. (2005). The Computational Theory of Mind, *The Stanford Encyclopedia of Philosophy*, (Fall 2005 Edition), Edward N. Zalta (ed.), http://plato.stanford.edu/archives/fall2005/entries/computational-mind.

[3] Ladyman, J. (1998). What is Structural Realism? *Studies in History and Philosophy of Science*, 29, pp409-424.

[4] Lycan, W.G. (1994). Functionalism, in *A Companion to the Philosophy of the Mind*, pp317-323, Samuel Guttenplan (ed.). Oxford: Blackwell.

[5] McGinn, C. (2004). What is it like not to be a brain?, in *Explanations: Styles of Explanation in Science*, John Cornwell (ed.). Oxford: Oxford University Press.

[6] Nagel, T. (1986). What is it like to be a bat?, reprinted in *The view from nowhere*. Oxford: Oxford University Press.

[7] Pooley, O. (2005). Points, particles, and structural realism, to appear in *Structural Foundations of Quantum Gravity*, S. French, D. Rickles and J. Saatsi (eds.). Oxford: Oxford University Press.

[8] Tegmark, M. (1998). Is 'the theory of everything' merely the ultimate ensemble theory?, *Annals of Physics*, 270, pp1-51. arXiv:gr-qc/9704009.

[9] Van Gulick, R. (2004). Consciousness, *The Stanford Encyclopedia of Philosophy*, (Fall 2004 Edition), Edward N. Zalta (ed.), http://plato.stanford.edu/archives/fall2004/entries/consciousness/.

[10] Votsis, I. (2003). Is structure not enough? *Philosophy of Science*, vol. 70(5), pp879-890.