

PHENOMENOLOGY & MIND

Edited by

Arkadiusz Chrudzinski · Wolfgang Huemer

Alexander Batthyany
Avshalom Elitzur (Eds.)

Mind and its Place in the World

Non-Reductionist Approaches
to the Ontology of Consciousness



PHENOMENOLOGY & MIND

Herausgegeben von / Edited by

Arkadiusz Chrudzimski • Wolfgang Huemer

Band 7 / Volume 7

Alexander Batthyany / Avshalom Elitzur (Eds.)

Mind and its Place in the World

Non-Reductionist Approaches
to the Ontology of Consciousness



ontos

verlag

Frankfurt | Paris | Enikon | Lancaster | New Brunswick

Bibliographic information published by Die Deutsche Bibliothek
Die Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliographie;
detailed bibliographic data is available in the Internet at <http://dnb.ddb.de>



North and South America by
Transaction Books
Rutgers University
Piscataway, NJ 08854-8042
trans@transactionpub.com



United Kingdom, Ire, Iceland, Turkey, Malta, Portugal by
Gazelle Books Services Limited
White Cross Mills
Hightown
LANCASTER, LA1 4XS
sales@gazellebooks.co.uk

©2006 ontos verlag
P.O. Box 15 41, D-63133 Heusenstamm nr Frankfurt
www.ontosverlag.com

ISBN 3-937202-98-6

2006

No part of this book may be reproduced, stored in retrieval systems or transmitted
in any form or by any means, electronic, mechanical, photocopying, microfilming, recording or otherwise
without written permission from the Publisher, with the exception of any material supplied specifically for the
purpose of being entered and executed on a computer system, for exclusive use of the purchaser of the work

Printed on acid-free paper
ISO-Norm 9706

This hardcover binding meets the International Library standard

Printed in Germany
by buch bücher dd ag

Table of Contents

Mind and Its Place in the World: Introduction and Overview <i>Alexander Batthyany</i>	7
What's the Mind-Body With You Anyway? Prolegomena to any Scientific Discussion of Consciousness <i>Avshalom Elitzur</i>	15
Thinkways: The Impulse to Reductionism <i>Hoyt Edge</i>	23
Self-Appropriation: The Dynamic Structure of Human Consciousness <i>Donald P. Merrifield, S.J.</i>	41
One Man's Meat is another Man's Person <i>Peter J. King</i>	61
Consciousness and the Intentional Awareness of Instantiatives <i>Russell Pannier and Thomas D. Sullivan</i>	77
Mental Monism Considered as a Solution to the Mind-Body problem <i>Peter B. Lloyd</i>	101
Telepathy: Or, How do I Know that this Thought is Mine? <i>Fiona Steinkamp</i>	145
The Dimensions of Conscious Experience: A Quantitative Phenomenology <i>Steven Lehar</i>	167
A Radical Externalist Approach to Consciousness: The Enlarged Mind <i>Riccardo Manzotti</i>	197
On Explanation, Interpretation, and Natural Science with Reference to Freud, Ricoeur, and Von Wright <i>Paul Lowland</i>	225
Personal Identity, the Self and Time <i>Howard Robinson</i>	245

Quantum Monism: Spinozicism Revived? 269
Gershon Kurizki

Boundary Conditions for Theories of Consciousness: 279
The Near-Death Experience and the Failure of Materialism
J. Kenneth Arnette

Contributors and Editors 317

Mental Monism Considered as a Solution to the Mind-Body Problem

Peter B. Lloyd

Consciousness has resisted conventional scientific attempts to explain it and, in the vacuum left by the absence of a scientific theory of consciousness, several metaphysical theories are in circulation. They fall into the three broad categories of physical monism, mental monism, and mental-physical dualism. This paper examines the second of these: mental monism (also called 'subjective idealism'), which is the theory that the conscious mental world is the primary reality, and that the physical world is a construct derived from it. The eighteenth-century philosopher George Berkeley was its clearest and best-known advocate in the West (Berkeley 1710). As this theory is diametrically opposed to the orthodox worldview of modern Western philosophy and science, it is appropriate to begin this paper by justifying even looking at this theory at all.

I will differentiate this theory of mental monism from a more popular theory that sometimes passes under a similar description. Certain interpretations of quantum measurement (following von Neumann 1932) assign a key role to 'consciousness', although quantum mechanics itself makes no mention of consciousness. In particular, some physicists (e.g. Wolf 1984, Goswami 1993) assign so radical a role to 'consciousness' that they claim to be idealists. But their concept of 'consciousness' is essentially that of quantum observation. In contrast, I am using the concept of 'consciousness' as it is employed by Chalmers (1996b) – essentially that of first-person qualitative experience. I assume that androids that lacked Chalmers-type consciousness could nonetheless carry out quantum measurements and therefore possess Goswami-type consciousness. Therefore the theory that I present here is unrelated to what might be called 'quantum consciousness'. Additionally, I

would differentiate the advocates of 'quantum consciousness' from others such as Henry (2005) who advocate mental monism in Chalmers' sense of consciousness but who nonetheless see quantum mechanics as a psychological tool in undermining a naïve belief in the external reality of a material world.

1 Respectability Arguments

1.1 Logical Respectability

Although the mind-body problem has received more intensive philosophical and scientific attention recently, after being renamed and made popular by David Chalmers as the 'Hard Problem of consciousness' (Chalmers 1996a, p 5; 1996b, p xiii), a generally agreed solution remains out of sight. Such intractable problems are often badly formulated problems, which stem from incorrect assumptions. So, let us examine the central assumption of the mind-body problem and ask: (a) what would become of the problem if we were to relax that assumption? and (b) what grounds are there for believing the assumption? If relaxing the assumption will solve the problem, and if the grounds for making the assumption are not logically compelling, then it is legitimate to investigate that approach. (What I mean by a 'logically compelling' assumption is one whose denial would involve self-contradiction.) Moreover, for the sake of completeness, we are obliged to investigate it.

Now, the key assumption in the mind-body problem is this: we have bodies, and our minds are dependent for their existence on our bodies – and the problem is to understand that dependence. For this exercise, therefore, let us suppose that we do not, in fact, have bodies and that, in consequence, the conscious mind can exist by itself without any physical support. (a) What happens to the mind-body problem? Clearly, it dissolves. (b) Are the grounds for believing that we have bodies logically compelling? No, for it is logically possible that one's whole life is a dream. Therefore we have established a justification for investigating mental monism as an approach to the mind-body problem. Needless to say, this is not an argument for mental monism as such,

nor even for the plausibility of mental monism. It is an argument with an almost vanishingly small, but nonetheless necessary, objective: it merely shows us that mental monism cannot be dismissed out of hand, and we have an obligation to assess it.

1.2 Conceptual Respectability

You probably have a firm and lifelong belief in physical reality, among the background assumptions that you never seriously question. You probably believe that the real world is largely, or wholly, physical and that you have a material body with which you interact in that world. When you read a paper in an academic publication, you naturally bring all those background assumptions with you, and expect the author to have the same ones in her background. If the author abruptly makes assertions that contradict those background assumptions, it can be confusing. Therefore, I will clarify my departure from one of those normal background assumptions, by using some philosophical thought-experiments, extrapolating from familiar situations.

1.3 Extrapolations

Consider dreams that are experienced during sleep. You find yourself in a world similar to the waking world, a three-dimensional space, surrounded by objects and people, and kept on the ground by gravity. You have a body in this world + at least, a notional point in space where your dreamed perceptions of vision, hearing, and touch are located. The dreamed objects will exhibit resistance to movement when pushed against, and can support the weight of your dreamed body. For instance, you might have a dream in which you climb concrete steps, and feel their rough surface and feel your body weight rising step by step. Nevertheless, that space and all the objects and bodies in it do not exist. We may say that they are delusory. While you are dreaming, you are having real conscious experiences that you construe as being produced by a physical world in which you are immersed. Yet that construal is false. That world does not exist.

I will now suggest three ways in which we can stretch our normal idea of a dream. First, a dream typically lasts for less than an hour. But, as a philosophical exercise, we can imagine a dream lasting for much longer: for a day, or a month, or a lifetime. Second, a dream is often narratively inconsistent — a living room might suddenly become a swimming pool, for instance. Nevertheless, lucid dreaming can establish greater degrees of coherence. Again, as a philosophical exercise, we can imagine a dream as consistent and reliable as our waking experience, even if it spans a long time, as much as a lifetime. Third, it is not uncommon to have dreams within dreams, and to dream that one has woken up. We can therefore imagine a long dream in which the dreamer repeatedly dreams that she falls asleep, and dreams, and wakes up. Putting these three extrapolations together, we can imagine a dream that is indistinguishable from normal waking life. Therefore, it is logically permissible to consider the hypothesis, again as a formal philosophical exercise, that your whole waking life is a dream. I am not saying that this is at all likely, or that there is anything to be gained by believing your life is a dream. I am merely noting that it would be logically coherent to have this belief, and the implications of its being logically coherent are interesting.

In fact, we have nowadays a better analogy than dreaming. For, dreaming as we know it is linked to a material brain, and uses imagery that has been gathered during the previous day from normal bodily perceptions. A congenitally blind person does not have visual content in her dreams; and if you anaesthetise a sleeper then she no longer dreams. A better analogy is a virtual-reality computer system. This involves goggles displaying pictures of an imaginary three-dimensional world, and a device such as a joystick or pressure-sensitive gloves, for moving one's avatar inside the virtual world. The appearance of this virtual world is, of course, conveyed entirely through one's sense organs by means of goggles and earphones. As a philosophical exercise, however, we can extrapolate the technology into the future and consider systems that interface directly with the nervous system. This idea has, incidentally, received popular attention in the science-fiction films *The Matrix* (Wachowski & Wachowski 1999) and *eXistenZ* (Cronenberg 1999), the former being examined by Lloyd (2003). Likewise we can imagine a computer so

powerful that it can deliver a comprehensive manifold of artificial signals, creating in the percipient's mind a realistic simulacrum of ordinary life. Again, we arrive at the logical possibility that one's whole life could be a computer-generated virtual reality. Whereas a dream has certain connotations, such as that it takes place for short periods of sleep, a virtual reality has no such traditional associations.

Indeed, it is logically possible to imagine that you have no brain, as your 'brain' in this virtual reality is only a virtual object in the virtual world. It is a virtual brain. You have no access to information about what substrate, if any, your mind is embodied in. For instance, one can imagine a science-fiction scenario in which a human mind is created and sustained in a virtual reality entirely within the memory of an advanced computer, without any biological tissue. Taking this extrapolation to its furthest extreme, we can frame the hypothesis that there is no physical universe at all, and that one's mind is created and sustained in a virtual reality driven by a purely mental computer. This is, at least, a comprehensible hypothesis in so far as it does not involve self-contradiction. And this is mental monism.

The difficulties that stem from taking this perspective seriously are very considerable, but I shall argue that they are soluble. Some of the most prominent problems are misconceptions of it that come from the sheer novelty of this perspective. I shall address two of those misconceptions in the remainder of this preliminary part of the paper — namely that mental monism is pragmatically untenable, and that it can be refuted by reaching out and touching the solid physical world. First, though, we need to do the following preparatory work.

1.4 Hierarchical Language-Games

The word "real" has several meanings, which are used in different circumstances. Hence, an unqualified assertion that the mental world is 'real', whereas the physical world is 'unreal', is ambiguous. I propose to articulate these different senses of 'reality' by borrowing a term from Wittgenstein, and describing our use of language in terms of 'language-games'. This is a neces-

sary preliminary to making an unambiguous statement and defence of mental monism. The term 'language-game' was introduced by Wittgenstein (1953), to denote a rule-bound pattern of social activity in which linguistic utterances interact with a wider range of actions. His point in introducing the term was to move away from a naïve Lockean view of words as tags and towards a perspective on an integrated system of language and behaviour. So, a word is no longer considered as having a single fixed meaning but as having different roles in different language-games. Wittgenstein gives the analogy of the levers of a steam engine, which look and feel the same (because they are designed to be handled) but perform different functions. Likewise, words might share the same look-and-feel but play roles of different kinds. In this essay, the main burden of the concept of language-games is to differentiate the roles played by terms of two particular classes: 'physical' terms and 'mental' terms. For instance, the preposition 'in' has the same look-and-feel in such sentences as these: "I have some food in my stomach" and "I have a pain in my stomach". But the meaning of 'in' is fundamentally different: in the first case, it denotes the spatial containment of an object; in the other, it denotes a psychophysical correlation. Nevertheless, that difference is obscured by the common linguistic appearance, and we may be led into the naïve belief that, in this example, the pain is physically located in the stomach.

Let us consider a more interesting case. According to Tarski's famous definition of truth, a statement such as "I am sitting in a noodle bar" is true if I am sitting in the noodle bar. Now, consider, on the one hand, sitting in a noodle bar in the physical world and, on the other hand, sitting in a virtual noodle bar in a computer-generated virtual world (such as in the Wachowskis' film *The Matrix*). In both cases you can truthfully say, "I am sitting in a noodle bar". So, what meaning is actually conveyed by that statement, and what does it refer to? The truth-conditions of the two utterances are, in fact, mutually exclusive: if I am sitting in the physical noodle bar, then I am not wired into the virtual world and therefore I cannot be sitting in the virtual noodle bar; and vice versa. But the immediate truth-tests are identical: I look around and see the tables and chairs, and the bowl of noodles, and I smell and taste the noodles; if I telephone my friend and ask her to join me, she will

have sensory impressions consistent with mine. Therefore, when my friend asks "Are you in the noodle bar?" and I say "Yes", the information I actually convey is about my sensory impressions, not about the unobserved putative substrate. So, as a matter of empirical fact, the meaning that is intended to be conveyed by the utterance is not the set of truth-conditions but rather the set of truth-tests. This is a crucial point, which academic philosophy systematically neglects.

In this connection, we have two different language-games: in the physical language-game, statements refer to truth-tests within the physical world; in the other, they refer to truth-tests within the virtual world. The same sentence belongs to one or the other language-game according to whether the statement is uttered in the physical or virtual world. In both cases, the truth-tests are the same, but the tests are carried out in different situations. That, however, is not the bottom of the matter. For, if my friend and I are acquainted with both worlds and can move between them, then when I invite her to join me, she may ask, "Which noodle bar? The real one or the virtual one?" We see that the two language-games are not on an equal footing, as the virtual language-game is subordinate to the physical one. This subordination consists in the fact that the results of truth-tests of the virtual language-game are logically supported (to use Foster's term) by facts within the domain of the physical language-game; and not vice versa. The contents of the computer database define the truth or falsity of statements within the virtual language-game, but those contents are facts within the physical world. Thus my being in the virtual noodle bar is a fact in the virtual world, but it rests on data that have been populated in a physical computer; and those data are physical facts. Consequently, in the superordinate language-game (in this case, the physical one), the world denoted by the subordinate language-game (in this case, the virtual one) is a construct – or, to put it bluntly, a fiction.

Let us pause to recapitulate. First, a given sentence may entail fixed truth-tests but also entail multiple different truth-conditions that depend on which language-game it is uttered in. Second, if one language-game is subordinate to another (in the above sense) then the truth-conditions of statements in the subordinate language-game are constructs or fictions that are logically

supported by facts in the more fundamental language-game. This conclusion is pivotal to removing a common misconception – namely, that mental monism is pragmatically refutable since it requires disbelief in the necessities of everyday life. How, for example, can I write this essay on mental monism if I believe that this laptop does not exist? The answer is that (a) in the physical language-game, the laptop can be said to exist; but (b) in the mental language-game, the laptop can be said to be a fiction, and (c) the physical language-game is subordinate to the mental language-game. The ‘mental language-game’ is the phenomenological language in which we report our conscious experiences. The physical language-game is subordinate to the mental language-game in so far as the results of truth-tests of physical statements are supported by facts in the mental world. For example, a truth-test for the (physical) statement that there is a potential difference of 1.5 volts between the terminals of the battery in this laptop would involve reading a voltmeter; but my consciously seeing a reading of 1.5 volts is a fact in the mental world, which is articulated in the mental language-game. It is in precisely this sense that all physical truth-tests are supported by mental facts; and that the physical language-game is therefore subordinate to the mental language-game.

With that distinction of language-games in mind, we can now dispel the apparent central contradiction of Berkeleianism. According to mental monism, it is true that (p) the laptop on which I am writing is real, but on the other hand, it is also true that (q) the laptop on which I am writing is a fiction. The proviso is that the sentence (p) is taken in the physical language-game, and the sentence (q) is taken in the mental language-game. This is a distinction of precisely the same kind as the one we make when talking in virtual realities, which allows us to make the true statement (r) I am in the noodle bar (in the virtual language-game), and (s) I am not in the noodle bar but wired into a virtual reality (in the physical language-game). Separating the language-games obviates the self-contradiction that is imputed to Berkeley’s philosophy. Berkeley himself did not have Wittgenstein’s idea of a ‘language-game’, but he did employ the corresponding idea of ‘vulgar usage’ and ‘learned usage’. This was ignored by generations of critics who asserted that mental monism, whilst it may be logically irrefutable, is pragmatically unten-

able. For how can a mental monist buy the groceries if she denies the existence of both shops and groceries? Berkeley’s answer is that, in the ‘vulgar usage’ (or the ‘physical language-game’), it is correct to assert the existence of the shops, whilst in the ‘learned usage’ (or ‘mental language-game’) it is correct to deny their existence.

This is not relativism. The physical language-game is subordinate to the mental language-game, because the truth-tests for the physical language-game are mental facts that can be stated in the mental language-game. Hence we may say (according to mental monism), the shop does not ‘really’ exist, although we go about everyday life as if we believed it did.

1.5 Comprehensiveness

There is an historical anecdote in which Berkeley visited his friend Jonathan Swift, only to have the door slammed in his face and to be told that, if the good Bishop was right in his denying the existence of the door then it could not bar his entry into the house. This implies that Swift thought that Berkeley’s theory of immaterialism applied only to doors and not to human bodies, a notion whose illogicality is apparent as soon as it is stated. In fact, neither door nor body exist in the mental language-game whilst both exist in the physical language-game. If Berkeley were to make the volitional efforts that would conventionally be described as trying to walk through the door, then he would experience in his mind a characteristic pattern of resistance to his movements. Precisely an analogous situation would exist if, in a dream, Berkeley were to try to propel his dream body through a dream door. There is another historical anecdote relevant here — that Samuel Johnson, when he was asked for his response to Berkeley’s immaterialism, gave a nearby rock a hard kick and said “I refute it thus”. This implies that Johnson thought that tactile perceptions gave him a direct connection with the material substrate. In fact, as Berkeley took pains to assert, tactile sensations exist within the conscious mind just as much as any sensation in any other faculty. Both these instances illustrate misconceptions of Berkeley’s philosophy that stem from supposing that mental monism is less than comprehensive.

1.6 Motivational Respectability

Mental monism is widely dismissed on the grounds that physical science has been overwhelmingly successful in accounting for the observed workings of the world, and its track record promises to deliver a comprehensive explanatory account of consciousness in due course. This has been called 'promissory physicalism'. The most visible counter-argument against promissory physicalism is its failure so far to make any progress in the mind-body problem. Despite three centuries of philosophical effort since Descartes first clearly stated the mind-body problem, despite the gains in brain science over the past century, and despite the efforts of consciousness researchers over the past decade (as measured, say, by the *Journal of Consciousness Studies*), the mind-body problem has not moved out of the philosophical arena into the scientific arena. It is still a philosophical problem, stated in terms that lie outside the physico-scientific vocabulary. At the basic level of designating the raw data of consciousness, namely conscious experiences themselves, we do not have physically defined terms. This places the mind-body problem in a different order from other hard scientific questions such as the nature of life, how gravity works, or how quantum mechanics works. In other such problems, the raw data on which the problem rests – such as reproduction and respiration in living organisms – are defined physically. There is something there for the physico-scientific investigators to formulate hypotheses about. Even early theories of life, involving 'vital forces', started from physical data. Likewise, the observable facts of gravitational attraction, or quantum-mechanical phenomena, are defined in physical terms. In all those problems, it is plausible that the physico-scientific method will ultimately achieve a comprehensive explanation, or at least asymptotically approach a comprehensive explanation. That plausibility is absent from the mind-body problem because it is not even defined in physical terms. We cannot even make the first hypothesis and kick-start the cycle of hypothesis and experiment, as the problem lies outside what can be addressed in physico-scientific hypotheses. Conscious content is known by immediate private experience, and not by public physical observation where it could be referenced by the third-person language of physical

science. In short, promissory physicalism cannot deliver a solution because the mind-body problem is not a physically defined problem.

Before the scientific investigation of the mind-body relationship can begin, there is a preliminary problem that needs to be solved, and that problem is a philosophical one. Since the problem is therefore necessarily outside the scope of the physico-scientific approach, we are bound to examine alternative approaches. The one such approach that would at once solve the mind-body problem is that of mental monism. This, therefore, provides the motivation to examine mental monism.

1.7 Summary

I have argued that mental monism should not be dismissed out of hand. It should be assessed on its merits, and in fact we are obliged to examine it for the completeness of the foundations of consciousness studies. I have argued for the respectability of mental monism in the following regards. (a) It is logically respectable as it is not self-contradictory and, if true, would solve the mind-body problem. (b) It is conceptually respectable as the world envisaged in mental monism is a consistent extrapolation of such familiar things as dreams and virtual realities. (c) It is motivationally respectable, as the most popular alternative to it, namely physical monism, cannot genuinely address the problem as it is not defined in physical terms.

2 Argument for Mental Monism

Having established that we ought to assess mental monism, let us now do so. The principal argument in its favour was first presented by George Berkeley in 1710, and it is generally known as Berkeley's semantic argument. This argument can be summarised briefly, as follows. Terms that denote conscious experiences are grounded by private ostensive definition, and thereby acquire a real semantic reference. In contrast, terms that denote physical things are defined analytically, and thereby are limited to a formal meaning

within a closed linguistic system. In consequence, physical things cannot really exist but are only formal constructs; whereas the contents of conscious experience are indubitably real. The remainder of this section will unpack this tersely stated argument. An instinctive reaction from anyone who has a scientific background is to dismiss any such semantic argument as mere chicanery. Surely the mental monists cannot hope to make the solid, material world disappear by a verbal sleight of hand? The mental monists' answer is that the supposed solidity of the material world is itself an illusion, conjured up by a verbal sleight of hand that has been perpetrated for centuries by the materialists.

2.1 *Mental and Physical Language-Games*

For Berkeleianism to make sense, it must be built upon the demarcation between mental and physical language-games. This foundation needs to be laid out before we can build the super-structure of Berkeley's semantic argument. This framework is usually omitted by the few philosophers who defend mental monism, who consequently leave themselves open to the sort of ridicule by which Geoffrey Warnock (1953) damned Berkeley for a generation. The idea of language-games is somewhat artificial, in so far as our real-life use of language involves sliding seamlessly between different language-games, making utterances in more than one language-game at the same time, and being unclear about which language-game we are using. Making headway in the mind-body problem requires that we pick apart the tangle of language-games, with more rigour than would be needed in everyday life.

Let us consider some comparatively simple term that is commonly used to denote a conscious sensation. With apologies to visually impaired readers, I shall choose red. If you are unable to see red, please substitute some other sensation, such as the sound of a whistle. The word "red" is used in three basic ways:

- (a) to denote a class of conscious sensations;
- (b) to denote a section of the electromagnetic spectrum;

(c) to denote a propensity of physical objects to reflect light in that part of the spectrum.

It can also be used, and most of the time in everyday life is used, to denote all three indiscriminately. Nonetheless, each of the three basic uses is also employed on its own in specialised language-games. (a) Artists discussing the colours to be seen in a particular scene and asking themselves which paints can capture the appearance will use the word "red" to refer to just the conscious sensation of red. So will dream analysts describing the colours they saw in their dreams whilst asleep. (b) An astronomer talking about red shift uses the same word exclusively to denote the electromagnetic radiation, and his statements would have the same meaning if no conscious being ever saw red – just as he talks about ultraviolet light. (c) A process engineer discussing how much red paint a factory produces in a day would use the word "red" to refer only to the reflective properties of the paint. For him, the word would convey the same meaning in a universe with no conscious beings and (*per impossible*) no light. Language-games of type (a) I shall call 'mental', in contrast with those of types (b) and (c), which are 'physical'. Some philosophers dispute that mental language-games exist in our pre-philosophical use of language, but their existence is demonstrated by the above examples. All three of these language-games (a) to (c) are pre-philosophical, in the sense that they arose in everyday life, and that is where they serve a useful role in conveying ideas and information.

2.2 *Private Ostensive Definition of Mental Terms*

So, let us consider the word "red" as it is used in the mental language-games, that is, those of type (a). How do you know what the word "red" means in those contexts? The common-sense view, which I shall be defending, is that you associate meaning with terms that denote conscious experiences by (i) having the experience in question, and (ii) mentally associating the word with the experience, or with a class of such experiences. For example, if I go to a paint shop to find out what a colour such as 'Tapestry Red' is,

then I look at the colour sample (Colour Council No. 12) and make a mental note, which might be articulated as, “*That is Tapestry Red*”. This is a ‘private ostensive definition’. It is ‘private’ (as opposed to ‘public’) because it takes place inside my conscious mind and is not amenable to scrutiny by third parties. It is ‘ostensive’ (as opposed to ‘analytic’) because the colour in question is fixed by mentally attending to it – figuratively speaking, I am pointing to the colour experience with my mind. It is not necessary that you should make a literal declaration (“*That is Tapestry Red*”) or deliberately intend to memorise the colour. I might merely be aware that the term is being used to refer to the colour. What is crucial is that the colour sensation occurs in the conscious mind. Definitions of this type cannot be established for humanly imperceptible ‘colours’ such as ultraviolet. I cannot learn ‘ultraviolet’ as a subjective colour term; it denotes only a segment of the electromagnetic spectrum. The term ‘ultraviolet’ is not part of the mental language-game. It exists only in the physical language-game. I cannot remember when I learned the names of more common sensory experiences such as ‘red’ and ‘blue’, as those occasions are lost in early childhood, but I can conceive of no other route than that of private ostensive definition.

2.3 Wittgenstein’s Private-Language Argument

I have argued that the private ostensive definition of mental terms is the only way we can give meaning to mental terms. It is widely believed, however, that Wittgenstein refuted the existence of ‘private language’ in general, and hence private ostensive definition in particular. On Wittgenstein’s view, what I have called the ‘mental language-game’ may *seem* to involve reference to mental states, but that reference is deceptive. Although Wittgenstein’s Private-Language Argument (*Philosophical Investigations* (1953), §269 *et seq.*) is frequently cited, it is not so often stated or analysed, partly perhaps because Wittgenstein’s gnomic style of writing militates against the concise statement of his arguments. Let us do so now. In a nutshell, Wittgenstein’s argument is as follows. Although we may naïvely believe that a term such as “red” can refer to a conscious experience, in fact it cannot do so. For, language is a

wholly public institution and lacks the wherewithal to refer to completely private mental experiences. I may claim that the word “red” refers to one of my experiences, and you may claim that it refers to one of yours, but as we cannot scrutinise and compare each other’s respective experiences, so we cannot establish the word “red” as referring to something in common between our minds. Nor can you establish it as referring to something constant over time within your own mind, since you cannot compare side-by-side the experiences you have at different times. A term such as “red” that you use in everyday life *as if* to express a private conscious experience is actually used to express a public fact only. If you use the language competently, then it expresses only the fact that the colour you are physically looking at is generally known as “red”. (That identification also brings with it a package of associated facts, all of them publicly observable, such as the wavelength.)

Let us now unpack this argument by means of an illustration. Wittgenstein wrote (§271): “Imagine a person whose memory could not retain what the word ‘pain’ meant – so that he constantly called different things by that name – but nevertheless used the word in a way fitting in with the usual symptoms and presuppositions of pain”. This is a fanciful example, but it is logically possible. In fact, we can follow Wittgenstein’s line of thinking further and imagine a person who never has the same sensation twice in any faculty – be it pain, vision, audition, whatever. Every time she looks at the same colour sample, she has a different colour experience; she also has a different sound experience when air waves of a constant frequency impinge on her ears; and so on. To complete this Wittgensteinian thought-experiment, we can suppose that this person is wholly convinced that she is having the same experiences whenever she is given the same sensory stimulus. So, when she looks at the paint sample ‘Tapestry Red’ on Monday and again on Friday, she has two quite different colour experiences, but she has the memory feeling as if the experiences were identical. We might say that her memory is ‘playing a trick’ on her. It is telling her that the experiences are the same even though, in fact, they are not. As long as her memory plays this trick consistently, neither she nor anybody else would ever know that anything is amiss. Wittgenstein uses this thought-experiment to illustrate the decoupling of language from

mental experience. Terms, such as “red”, that nominally denote the contents of conscious experience, actually have no such meaning. Instead, they play a different role in the language-game. They are indices of publicly observable facts only. When I look at a colour sample and claim that I see the colour red, I am not (according to Wittgenstein) actually reporting a conscious colour experience, but am conveying only that this paint is what we normally call “red”, and by implication it will possess the normal publicly observable properties of red. For example, that other competent speakers will also call it “red”; and that a laboratory technician could measure its reflected light and find its wavelength to be between 740 and 620 nm. This is not to say that a person uttering “I can see red paint” is thinking about the wavelength of the light. Nevertheless he has identified the paint colour’s role in the language-game, which brings with it a package of implications including the publicly measurable wavelength.

Wittgenstein is not denying that we have conscious experiences. He is not arguing (as Gilbert Ryle did) that there is really no colour experience in the mind. Rather, he is arguing that our language is isolated from whatever mental experiences there may be. Wittgenstein bases this theory on how, in principle, utterances can acquire their meanings. The members of a linguistic community can establish the meaning of an utterance by ‘public ostensive definition’ (for example, by pointing to a sample of Colour Council No. 12 and saying, “This is Tapestry Red”) or by analytical definition (for example, by stating the spectral distribution of its reflectance). In practice, there are a myriad ways in which these definitions can be carried out. Most often, you just notice how other people use words, and you unconsciously interiorise that knowledge – which, in effect, is an informal ostensive definition. Private mental experiences, however, are not amenable to definitions of either type. Therefore, says Wittgenstein, no utterance can be established as meaning a private mental experience.

To counter Wittgenstein’s Private-Language Argument, I shall argue that his reasoning applies to all languages, not just private ones and therefore entails the *reductio ad absurdum* that all communication is impossible. This, in fact, is the burden of a point that A.J. Ayer made in his biography *Wittgen-*

stein (Ayer 1985, p 75). Suppose that all observable properties of external bodies were to be in continual flux, but that some mischievous intelligence (‘Wittgenstein’s demon’) were continually to recalibrate all measuring instruments and all sense organs, so that everybody still ended up reporting the same observations. For example, suppose the wavelength of light from a blue sky were to shift from one value to another in every moment, but Wittgenstein’s demon simultaneously recalibrated all spectrometers to give a constant reading. Then the physical term ‘blue light’ would not refer to any fixed thing. Likewise for every other observable characteristic. Thus our public language would carry on just as it does now, even in this imagined universally changing world. Therefore, on Wittgenstein’s reasoning, we must conclude that language never means anything. It is a rule-based social activity but signifies nothing. Wittgenstein’s Private-Language Argument fails because direct, public ostensive definition – which he assumes to underpin public language – is not feasible. Strictly speaking, there are no public facts, only facts of varying degrees of privacy. You cannot, in principle, check your perception of a public thing, such as a colour sample, because any checks you make depend on your sensory observations and could be tricked by Wittgenstein’s demon. Even if you ask other members of your linguistic community to confirm that you are using the word “red” competently, the demon may trick all of them into mistaken agreement. Language is at one remove from supposedly public facts. For those facts are not directly knowable, but only inferred from immediate mental experiences – which are themselves private and therefore outside the scope of public ostensive definition. In the public language, utterances are thus not coupled to their meanings in the direct way that Wittgenstein needed. That coupling would require both speaker and hearer to have unmediated access to the supposedly public facts that constitute the meaning of each utterance. In fact, they cannot have such access. In the case of private mental facts, only the speaker has unmediated access; in the case of public external facts, neither party has unmediated access.

In practice, we establish the public meaning of both mental and external terms *conditionally*, on the *assumption* that the external world is stable, and that the correlation between mind and brain is stable and universal. With

those assumptions, you can point to a sample of Tapestry Red and be confident that you and I will have very similar mental experiences. You can also be confident that, when you hear other people say it is 'Tapestry Red', they are not deceived, and you hear aright. By this means, human communities communicate subtle shades of emotion and sensation. So, for example, in artistic fields as diverse as music and whisky-tasting, private ostensive definitions are routinely communicated between members of those specialised linguistic communities. Thus humankind has built up a language in which we can competently discuss both private mental facts and public external facts. Admittedly, it all rests on the aforementioned assumptions, but, for all practical purposes, those assumptions have served us well for millennia. We have therefore rebutted Wittgenstein's Private-Language Argument as an objection to our notion of private ostensive definition. We can also reject Dennett's (1988) use of it. In his usual style of allusive argument, Dennett puts the reader in mind of Wittgenstein's Private-Language Argument in order to support his denial of qualia, although he does not explicitly infer anything from it. With Wittgenstein's Argument defeated, however, whatever support Dennett seeks to derive from it likewise fails.

2.4 Analytical Definition of Physical Terms

We have examined how meaning is assigned by private ostensive definition to mental terms. We will now examine the assignment of meaning to terms used in the physical sciences, which rest on analytical definition. To clarify 'analytical definition', let us first consider a trivial example from set theory. Consider a set $X = \{p, \{a, b, c\}\}$ including the operation $p:a \rightarrow b$, $p:b \rightarrow c$, $p:c \rightarrow c$. None of these terms has any meaning beyond the formal one that is explicitly contained in this definition. No amount of further scientific research can discover new properties of p . Nor does it make sense to ask about aspects of p that are not covered by the definition, such as what the mass or the volume of p are. There is an illuminating parallel between mathematical constructs, such as the one above, and literary fictions, such as *Sherlock Holmes*, and a comparable analogy with computer software, such as a program written

in Fortran. In each case, there is a text comprising statements about certain entities, which are exhaustively and comprehensively defined by what is written about them in that text. Just as the operation p can have no properties beyond those defined above, so also the person of Sherlock Holmes can possess no details beyond those written by Conan Doyle, nor can an integer variable J have any properties other than those declared in the Fortran program. It would be meaningless to ask whether Holmes, upon rising in the morning, put on his left shoe or his right shoe first, just as it would be meaningless to ask what colour J is. In just this sense, we may say that the entities a, b, c , and p are all fictions.

A fiction need not remain forever decoupled from the outside world. It can be bound to the world by explicitly associating each term in the text with something in the observed world. We do this by public ostensive definition: for instance, the fiction of a perfect circle might be applied to some piece of engineering in order to compute its geometry. (In physics, the text is a 'model' for the relevant part of the outside world. In literature, the text describes an 'allegory' of the world. In computer science we say that a programming language has a 'binding' in a particular computer.) Now, returning to our simple example, we can apply the set X as a model of, say, the coffee mug on my table. We can do this by defining an 'interpretation' function N that maps the terms into the relevant external set. So, let $N:a \rightarrow$ 'the mug is in the middle of the table'; $N:b \rightarrow$ 'the mug is near the edge of the table'; $N:c \rightarrow$ 'the mug is on the floor'; and $N:p \rightarrow$ 'I give the mug a push'. In this trivial example, the separation between the linguistic model and the modelled world is inescapably obvious. In more complex and interesting examples, the customary terminological slackness that mathematics is fond of may tend to blur this distinction. And, in examples in everyday language, the distinction is buried in layers of blur. Chess offers us another example of a fiction (namely the rules of chess) that has a binding in the physical world. For example, in a game of chess, the pieces *per se* are abstractions, which are defined wholly by their role in the game. But the game can be 'bound' to any individual physical chess set. Then you might pick up a wooden piece in your hand and loosely say "This is the King". Of course, the King remains a relationally defined ab-

straction, and the wooden King piece is playing that role, but we habitually talk as if the wooden piece were actually the King. Without adequately careful thought, we can easily be misled into conflating the abstraction with its token, and think that the tangible chess piece is the abstract one, that this piece of wood *is* the King. The philosophically relevant principle that we need to take away from these elementary cases is that the terms that denote fictional things still do so even when the fiction has a binding in some external world. The symbol “p” itself still denotes a fiction, even when that symbol has a binding to my pushing my mug; and in chess “King” still denotes a fiction even when it has a binding to a particular piece of wood. It is not too difficult to bear this distinction in mind when dealing with such shallow fictions. What is harder, but crucially important, is to continue to bear it mind when we deal with deeper fictions.

There is a contrary claim that we have to consider here, which I shall call ‘anti-fictivism’. This is the claim that when a formalism such as X is bound to an external world, it ceases to be a *fiction* and becomes a *description* of the mapped part of the external world. This claim does not go through because of the ‘topic-neutrality’ of the fiction. Topic-neutrality is a term introduced in this context by Gilbert Ryle (1949) and developed by J.J.C. Smart (1963). It means that the fiction picks out certain logical features of the denotand but leaves completely open the intrinsic character of what is denoted by its terms. In a nutshell, the argument is as follows. If an account, such as X, is topic-neutral, then it leaves undefined the intrinsic character of what it denotes; but any really existing thing will have an intrinsic character of some sort; therefore an account such as X cannot be about any really existing thing. That is to say, it is fictive.

A corollary of topic-neutrality is that a fiction can have many different bindings. My trivial formalism, X, can serve as a model not only of my mug, but also of any other object lying on the table, such as my pen. Another example is that of the formalism of chess, which applies to every chess-set in the world, whatever size, colour, or material. The topic-neutrality of shallow fictions such as these is obvious. I have dwelt upon it in order to make the principle stand out. For, the same principle applies to the topic-neutrality of

the deeper fictions of physics.

Before we move on to consider physics, there is a further consideration I would offer to bolster my position. Returning again to the formalism of X, let us now suppose that I have a virtual-reality program running on my computer, which displays a graphical image of my mug on my table. Instead of actually pushing the mug, I can press the arrow key on the keyboard, which gives a virtual push to the virtual mug, which in turn is portrayed on the screen. Now re-bind the formalism of X, so that $N:a \rightarrow$ ‘the virtual mug is in the middle of the virtual table’, $N:p \rightarrow$ ‘I give the virtual mug a virtual push’, and so on. There is no longer any attraction to claiming that the formalism of X is a *description* of the external world, rather than a *fiction* that happens to be bound in a certain way. For, the codomain of the interpretation function, N, does not strictly exist. Since the virtual mug and the table do not actually exist but are only simulations on a computer, it cannot be contended that X is describing something in the external world. Here, the fictivity of X is transparently obvious, but my argument is that this X is always a fiction because it is topic-neutral.

The other philosophically significant point to take away from this simple example is that what is observed (namely, the coloured pixels on the screen) is orthogonal to what the formalism X denotes (the mug and the table). In this particular example, an inscrutable mechanism – the computer software – operates a virtual reality that manifests an observable world structured as a virtual three-dimensional space that is depicted on the screen, obeying a formalism, or set of mechanical laws. According to Berkeley, a similar schema applies to our everyday world. The ‘inscrutable mechanism’ is God, and the ‘structure of the observable world’ is all that the physical sciences can find.

2.5 The Fictional Nature of Physical Things

Topic-neutrality is a general concept and is not limited to the particular kind of topic-neutrality that we have been considering. Ryle and Smart originally ascribed topic-neutrality to consciousness, but here we follow Foster’s

(1982) opposite conclusion of ascribing it to physics. Foster notes the irony of this reversal thus: "It seems to me that this way of defending physicalistic realism against the intuitive objection fails. It seems to me that, in a sense, it fails catastrophically, since it manages to get everything exactly the wrong way round. It locates topic-neutrality at the very point where our concepts and descriptions are topic-specific, and it locates topic-specificity at the very point where they are topic-neutral." As we shall now see, it is physics that is topic-neutral, not the phenomenology of the mind. Physics posits a number of entities, traditionally the proton, neutron, and electron, but now encompassing a zoo of elementary particles. Each particle is exhaustively and exclusively defined by its mathematical properties – its mass, charge, spin, and so on. On this view, the proton (like all the other particles) is, strictly speaking, a mathematical construct, or fiction. The proton is topic-neutral precisely because it is exclusively and exhaustively defined by its extrinsic relations and not by any intrinsic qualities. It is nothing but a construct that is wholly defined by that combination of extrinsic mathematical properties.

There is a key objection that every physicist will raise here: "Surely, the proton is a real thing that exists out there in the external world, and the mathematical properties are a description of the proton, not a definition of it". That, however, is precisely to forget what we learned above from considering such simple examples as set X, and the game of chess.

The account of the proton and other particles is clearly topic-neutral, because the theories of physics allude to only the properties of the particles that are enumerated in their mathematical definitions – their mass, charge, spin, and so on. The collective text of physics gives us a 'model', or as I would like to say, a 'fiction'. That model can be given a binding in the structure of the observable world, but because the text is topic-neutral, the terms of the collective text of physics still denote constructs, or 'fictions'. Furthermore, the entities that are denoted in the text of physics are inherently inscrutable because an observer can directly observe only the contents of her conscious perception. It cannot seriously be contended that anybody has direct conscious experience of protons and other elementary particles. Nor can any seriously reflective person propose that we can make direct observations of

macroscopic bodies that are supposedly made up of elementary particles. Lloyd (1999a, section 2.2.4) argues at length against such 'naïve realism'. This analysis is akin to that of Margenau (1984) who emphasises that what I have called the 'structure' of the observable world (his 'constructs') exhibits a gradation of abstraction away from what is consciously observed.

So, the conclusion that we are drawn to is that the physical world is a fiction. The resistance to this conclusion, especially among scientists, cannot be overestimated. Vehement denial is the normal response, but evoking a coherent counter-argument is not so easy. So, instead, we have to analyse some of the more visceral responses to it.

The responses can largely be characterised as the claim that we 'know' that physical particles are not fictions, because we can reach out and touch tangible macroscopic objects, such as a table, which are assemblages of those very particles. This is a conflation of immediate perception with abstract constructs that are projected onto it. When you touch the table, you have a conscious tactile experience: this is *inside* your conscious mind. It is not a bridge from the mental to the physical, it is not a trans-mental encounter with an assemblage of atoms: it is an event that occurs within the realm of consciousness. Onto that conscious experience is mapped the physical event of your fingers touching the surface of the table (the first level of Margenau's constructs), and from that you infer that the atoms in your finger are encountering the atoms of the table (a deeper level of Margenau's constructs). Nevertheless, the feeling of solidity, of resistance to movement, of the brute fact of 'being there' – is all, literally, in your mind. It is a pattern of conscious experience. And that mental experience of resistance is precisely what, I believe, people are alluding to when they say that they know that matter is really 'out there'. They are saying, in effect, that there is a manifold of tactile resistance that the conscious mind engages with. But that conscious experience of solidity is conceptually quite separate from the mathematical model of the physical system. To object that the physical system cannot be a fiction because you can feel it to be solid is to lose sight of that conceptual distinction between perceived solidity and a mathematical model.

2.6 Conclusions: Part One

From the earlier sections, we must conclude that what is denoted by mental terms is directly observed, for otherwise it could not be denoted by mental terms – and therefore it must be objectively real; while what is denoted by physical terms is an abstraction by its definition, and hence a construct. From this, the conclusion of mental monism is inescapable. At the start of this paper, we sought to assess mental monism. We saw that it was a logically possible solution to an otherwise seemingly insoluble problem, but we recognised that it was universally derided as absurd. Nonetheless, for the sake of completeness, we were obliged to examine the argument for mental monism. Now, however, we see that the twin pillars of Berkeley's semantic argument for mental monism are sound: on the one hand, the veridicality of the phenomenological language; and on the other hand the topic-neutrality, and hence fictivity, of physics. Thus mental monism is supported by this argument. Of course, this is an unwelcome conclusion. Therefore we shall look further afield, for some other theory that can accommodate that argument but save us from mental monism.

2.7 Neutral Monism

When confronted by the realisation that the mental world is not reducible to the physical, a common move for physicalists is to migrate to neutral monism. This is a theory originally associated with Bertrand Russell (1927), which has gained followers in recent years, as research in consciousness studies reveals the inadequacy of physical monism. Its central claim is that there is some ultimate substance of the universe that is neither physical nor mental. Chalmers (1996b, p127) refers to such stuff as having 'protophenomenal' properties. How does neutral substance differ from physical substance? An account of neutral substance is topic-neutral, while in contrast physical substance is supposed to have some uniquely physical intrinsic character. As we have seen above, however, any account of *physical* substance is also topic-neutral. The enterprise of neutral monism collapses as soon as it is recognised

that physics is, in fact, already topic-neutral and there is no intrinsic physical character. More precisely, my argument against neutral monism is that, (i) any terms that do not denote mental things must be defined analytically rather than by private ostensive definition, and therefore must be topic-neutral; but (ii) physical terms are topic-neutral and therefore denote no specifically physical intrinsic character, and hence (iii) protophenomenal stuff is, in fact, no different from physical stuff – and neutral monism is, in effect, a variant of physical monism. Therefore we can conclude that the neutral monistic substance, the protophenomenal stuff, is a fiction (or construct) just as physical substance is. Another way of looking at it is that the 'Hard Problem' resurfaces in neutral monism. Let N be the neutral-monistic facts associated with my seeing this table. Why is N accompanied by conscious experience? In its essence, this problem is the same as the original one: let P be the *physical* brain facts associated with seeing this table; why is P accompanied by conscious experience? This sterility of neutral monism leads its proponents to smuggle consciousness in by the back door, giving rise to the theory that I shall call 'mental-structural dualism'.

2.8 Mental-Structural Dualism

René Descartes was the first and most famous proponent of full-blooded dualism (Descartes, 1637). According to that theory, there are characteristic intrinsic properties of both physical things and mental things, which are mutually exclusive. Hence Descartes was a 'substance dualist': he said there were distinct physical things and mental things. This theory could not explain interactions between physical things and mental things. For example, if minds are not located anywhere in physical space, why should my mind interact with *my* brain rather than yours or anybody else's? Trying to tackle that problem led to 'property dualism', in which there was supposed to be only one set of things, but they had properties of two kinds, mental and physical. This, of course, contradicted the non-spatial nature of the mind — which is what had driven Descartes to dualism in the first place. Both of those forms of full-blooded dualism inevitably fail because, as we have seen above, there is

no characteristic intrinsic quality of being physical. Physical terms are topic-neutral. So the physical half of any such dualism collapses to vacuity. The next theory is anaemic dualism, or what I shall call 'mental-structural dualism'. This theory maintains that the constituents of the world have intrinsic phenomenal qualities, but are structured as physical things. Lockwood (1987) expressed this by saying that the conscious mind is the physical brain seen from the inside. Rosenberg (2004, p 78) briefly outlines the history of what I call 'mental-structural dualism' and he calls 'liberal naturalism'. Mental-structural dualism is a half-way house between neutral monism and property dualism. Neutral monism says that ultimate reality is neutral in the sense that it contains neither physical nor mental intrinsic qualities. Property dualism says that ultimate reality is both mental and physical, and contains intrinsic qualities of both sorts. Between those two theories, mental-structural dualism allows that reality possesses intrinsic mental characteristics, together with structural and extrinsic physical characteristics – but not any *intrinsic* physical characteristics.

Eddington (1928) advanced a naïve early formulation of what I am calling mental-structural dualism: "We realise now that science has nothing to say as to the intrinsic nature of the atom. Our knowledge of the objects treated in physics consists solely of a schedule of pointer readings. The schedule is we agree attached to some unknown background. Why not then attach it to something of a spiritual nature of which a prominent characteristic is thought". Lockwood (1998, p 90) advocated a specific form of mental-structural dualism: "My proposal is that a quale is the intrinsic reality corresponding to an occupied point within an n-dimensional quality space, points within which are defined by n-tuples of simultaneously realisable eigenvalues of a set of n compatible observables on the relevant subsystem of the brain."

In his doctoral thesis and subsequent book, *A Place for Consciousness*, Rosenberg (1997, 1999, 2004) has elaborated a version of mental-structural dualism, to which he gives the oxymoronic name of 'pan-experientialist neutral monism'. (It is oxymoronic in so far as 'neutral monism' posits a reality devoid of conscious experience, contrary to 'pan-experientialism'. It is also misleading, as he advocates, not neutral monism, but what I have called men-

tal-structural dualism.) Discussing parallels between his work and that of Lockwood (1989, 1993), he writes, "Lockwood suggests that phenomenal consciousness fails to logically supervene on the physical because physical concepts are content-neutral, merely specifying the structure of the causal flux. Phenomenal qualities and consciousness, on the other hand, are defined precisely by their content" (Rosenberg 2004, p 111) – which is a neat summary of why physical monism cannot succeed. He continues, "Lockwood suggests that a nice solution is simply to draft phenomenal properties into duty as the content of the causal flux whose structure is described by physics" — which is a neat summary of the mental-structural dualist theory that results from introducing consciousness into neutral monism, as Rosenberg himself does. Rosenberg, along with Lockwood and Chalmers before him, want to bring consciousness in at the basement level, as a fundamental element of reality. All of them seek to achieve this by proposing that there is phenomenal content of some sort associated with even the most finely grained physical process.

The individuation of human minds is a challenge to mental-structural dualism. If every elementary particle has some minuscule phenomenal property, how does the phenomenal world consist of private streams of consciousness corresponding to individual human brains? We should have expected either that each elementary particle would have mental privacy or that there would be a single universal mind. Rosenberg's objection to Lockwood is primarily that Lockwood does not give a principled account of this individuation of the phenomenal world into personal minds, whereas Rosenberg gives a very complicated theory to try to explain it. Otherwise, the philosophies are very close. Even with his account of the individuation of human minds, Rosenberg has to bite the bullet of panexperientialism. Any plausible theory of mental-structural dualism (or, indeed, property dualism) that ascribes phenomenal properties at the microlevel entails that conscious experience pervades every nook and cranny of the universe. This does not faze him, though: "it is not even clear in what sense the intuition against panexperientialism really is a common sense intuition. Many other cultures have seriously entertained or endorsed an animistic metaphysics and it is certainly possible that

the current rejection of distantly related views like panexperientialism is somewhat of a knee-jerk reaction against what are seen to be more primitive or theistic views.” Likewise, Strawson (1999, 2005) proposes a mental-structural dualism with pan-experientialism at the micro-level. He has, however, made interesting claims about its relation to the monisms, which we shall now look at.

2.9 *Mental-Structural Dualism versus Mental and Physical Monisms*

Mental-structural dualism is a special case of *mental* monism, in the following respect. Mental monism’s defining tenet is that reality is fundamentally mental, whilst mental-structural dualism’s defining tenet is that reality is fundamentally mental, but structured by the ontology of physics and subject to the nomological constraints of physical laws. Strawson (1999, 2005) makes the opposite claim that mental-structural dualism is a form of *physical* monism. This has been misread in some quarters as the merely terminological tactic of redefining the word ‘physical’ to encompass the mental. In fact, Strawson’s starting point is the topic neutrality of physics, which leaves open the intrinsic nature of the physical, and he proposes that the intrinsic nature of the physical is mental. In this way, he is able to assert that reality is just the physical world, but it includes the mental world as its intrinsic nature. Strawson argues validly but from the false premise that the phrase ‘physical world’ *refers* to an external reality, which is *described* by the laws of physics. On the view that I defended earlier, however, the physical world is a *model* or construct which we *bind* or project onto the regularities of the phenomenal world. Therefore, the phrase ‘physical world’ does not bear a semantic reference to a real thing, and the laws of physics define a closed formal system and do not describe an external reality. Physical monism asserts that the physical world is all that exists, but that cannot be true because it ascribes an independent reality to something that is purely notional. Trying to ascribe phenomenal qualities to physical things is as incoherent as trying to ascribe them to any abstraction such as the number five. In contrast, mental-physical dualism is at least coherent, as it ascribes extrinsic physical relations to really ex-

isting conscious phenomena. Nevertheless, Strawson’s assertion of physical monism – that is, that reality consists of the physical world – would still be a useful metaphor *if* mental-physical dualism were true – that is, if the phenomenal world were actually isomorphic to the physical.

Can mental-structural dualism be true? No, because the binding of micro-experiencers into macro-experiencers imposes an untenable spatial ordering on the mental world. As I argued at length elsewhere (Lloyd 1999a, section 2.4), if spatial location could be predicated of mental things (as it can be of physical things) then a conscious mind could be subjected to spatiotemporal dissection, but that contradicts the observable unity of the conscious mind. Therefore, by *reductio ad absurdum*, consciousness cannot be spatial. In a nutshell, the argument given by Lloyd (1999a) is as follows. Suppose that two mental facts M_a and M_b are co-conscious, i.e. they occur at the same psychological moment in your mind, but they are identified with spatially separated brain events B_a and B_b . Suppose that the physical correlate of the co-consciousness of M_a and M_b is the transmission of some causal influence C between B_a and B_b . Now, conduct an experiment in which B_a and B_b occur simultaneously at time T_1 but C is delayed substantially until T_2 . Then mental-structural dualism predicts that M_a and M_b will occur at T_1 but will not be co-conscious until T_2 . This contradicts the psychologically necessary fact that M_a and M_b must be co-conscious when they occur and cannot become co-conscious retrospectively.

Despite Rosenberg’s theory of how personal minds are individuated, mental-structural dualism cannot avoid the assumption that minds are extended in space. But, by the dissection argument outlined above, any theory of spatially extended consciousness must fail. We are thus forced to relinquish mental-structural dualism in favour of a Berkeleian mental monism, which takes macro-experiencers as primitive. Hence also, Strawson’s assertion of physical monism is not valid even as a heuristic metaphor.

2.10 *The Metamind*

Like Berkeley in his *De Motu*, Rosenberg correctly recognises the hid-

den force of causality as belying the active presence of an agency that ultimately makes things happen. Whilst the laws of physics may prescribe the form of what can happen, they necessarily fall silent on what makes the real world toe the nomological line that physics measures out. Rosenberg's answer is that consciousness is the *carrier* of that agency. But what is the *agency*? Berkeley's answer was that the agency is the will of God. This answer is empirical, in so far as we have an acquaintance with volition in our everyday lives, as manifested in the free will of our voluntary acts. But does the answer have any explanatory power? Before we discuss this, I shall propose a change of nomenclature. The word "God" has a vast, and culturally charged, hinterland of connotation and association. Whilst leaving aside the question of whether Berkeley was right, I propose that we use the term 'metamind' (Lloyd 1999a) to denote the posited mind that drives all natural phenomena that we ourselves do not will.

In the Berkeleian picture, therefore, the world comprises phenomenal qualities, grouped and associated somehow, in which there is a single, uniform kind of agency — volition — with which we have direct acquaintance in everyday life. All processes that are not driven by our own volitional acts are driven by volitional acts of another mind, the 'metamind'. Berkeley largely glossed over the internal structure of the metamind. He asserted that the metamind contained within itself 'archetypes' of the macroscopic objects that we see around us as the furniture of the world. All he would say of those archetypes was the negative observation that there was no reason to think that they resembled what we experience mentally when we observed the corresponding objects. His analogy was that of a musical score, which does not resemble the music that the player renders it as in our hearing. Likewise, God's archetypes need not resemble the perceptions that God renders them as in our minds. That Berkeley did not attempt to penetrate the metamind does not preclude our attempting to do so. On the contrary, the scientific imperative dictates that we seek to analyse and understand in full the workings of the metamind. We must, as it were, reverse engineer the manifest world in order to discover the informatic structure of the metamind. Only then will the metamind serve as an explanatory concept. This, then, measures the distance from

mental-structural dualism to mental monism: mental-structural dualism presupposes (mistakenly, as I have argued above) that the metamind is isomorphic to the physical world; whereas mental monism denies that the metamind can have a physicalistic structure but regards the actual structure as a subject for scientific investigation and discovery.

There is, I think, an illuminating parallel to be drawn with virtual reality systems. If we compare our manifest phenomenal world to the visual, auditory, and tactile imagery that is output by a virtual-reality computer system, then the metamind corresponds to the software driving the simulation. Within that software, there are modules that simulate individual objects, and which can render the appropriate imagery for any observer's particular point-of-view. Once we start thinking of the metamind in this way, we find a potentially fruitful framework for formulating hypotheses about the informatic structure of it. We have a point of departure for probing the metamind.

2.11 Conclusions: Part Two

As we have seen, the collective discourse of physics is topic-neutral. It defines logico-mathematical structures but fails to make semantic reference to the basic reality of the universe. In contrast, the collective discourse of consciousness succeeds in making reference to an immediately apprehended world of phenomenal content. Physics defines a model of the observed structural patterns of the phenomenal world. The physical model cannot be isomorphic to the mental world, because the mental world exhibits unity at the level of private personal minds. The processes of the mental world that are not driven by personal volition are attributed to the 'metamind', whose internal structure must be discovered by scientific investigation.

3 Modelling the Mind-Brain Interface

Although the mental and physical worlds are ontologically different, we know from every moment of waking life that there is a close correlation between at least part of the mind's contents and part of the physical world. You have sensory experiences that correlate with the patterns of energy transduced by your body's sense organs. In short, you see what your eyes look at, and so on. Furthermore, it is well known from elementary biology that the true locus of correlation is not at the outward-facing surface of the sense organs but at some internal points of cognitive processing in the brain. There are two important questions here. (a) Where is the locus of the mind-brain correlation? (b) What does the correlation signify?

3.1 Locus of the mind-brain interface

By 'locus' I do not mean just which anatomical structure it is in, but rather what kind of structure or aspect of brain tissue it is. A key clue lies in the fact that the conscious mind can affect the brain. This is demonstrated by the fact that your body can report your conscious mental states. Yet the physical world is subject to its own internal nomological constraints: there are causal laws that, for a given physical state at time T_1 , limit what the physical state can be at a later time, T_2 . In a deterministic universe (for example, one obeying classical mechanics, such as a Newtonian universe), with fixed initial conditions, the nomological constraints are closed in the sense that the total state of the universe at any time is predetermined. Such a universe could not manifest conscious minds. A brain could never report conscious states in such a universe, because the physical state at each successive moment in time is determined entirely by the antecedent physical state irrespective of anyone's conscious mental content.

Therefore, physical non-determinism is a requisite for the mind-brain interface. And the locus of the mind-brain correlation must comprise non-deterministic events or conditions. There are two possible kinds of non-determinism in this context that could serve as an interface through which con-

sciousness could manifest in a physical world. Namely, the initial conditions, and causal discontinuities. We shall examine each in turn.

3.2 Open Initial Conditions

A deterministic (e.g. Newtonian) universe with fixed initial conditions cannot form an interface with consciousness. But if some unobserved initial conditions are not fixed, and if the world possesses dynamical systems that are mathematically 'chaotic' then there is, in principle, an opportunity for an interface into consciousness. (A system is mathematically 'chaotic' if its macroscopic behaviour is sensitive to arbitrarily small changes in its antecedent conditions. No matter how precisely you measure the system's state at time T_1 , the residual measurement error will be sufficient to hide the early causes of macroscopic effects at subsequent time T_2 .) To articulate how the initial conditions can offer an interface to consciousness, it is crucial to differentiate the physical time and mental time. The former is just the conventional linear dimension in which events are described as occurring in the textbooks of physics. The latter dimension is defined by the succession of mental events.

Recall that, by the arguments of the preceding section of this essay, the physical world is a construct invented by the conscious mind to rationalise the observed structure of the phenomenal world. It does not have a mind-independent reality. Therefore the physical universe as a whole, including its entire course of development over physical time, has a dependence on mental time. Let $U(T_p)$ be the state of the physical universe at physical time T_p , and let $W = \{U(T_p): \text{for all } T_p \text{ up to now}\}$ be the ordered set of all such states until now. Now, the natural view is that $U(T_p)$ is a function of mental time T_m , and we could write $(U(T_p))(T_m)$ to express its dependence. What this means is that the community of conscious minds (or, for that matter, an individual conscious mind) asserts a body of physical statements that, taken together, defines a physical universe. Needless to say, what we choose to define as the physical universe is driven by the observed patterns and structures of the phenomenal world, which is in turn driven by the metamind. So there is a deeper dependence on the metamind. But the metamind, in so far as it is a mind,

evolves in mental time, and so it is encompassed by the term T_m in the representation $(U(T_p))(T_m)$.

What I would propose, however, is that not only the current universe U but the whole time-series $W = \{U(T_p)\}$ depends on mental time, thus: $W(T_m)$. Stated bluntly, this means that the physical past can be changed in the mental present. In order for this to be a realistic model, there must be a constraint of backwards compatibility. It would be absurd to propose that things that were observed to happen in physical-yesterday should retrospectively, in mental-today cease to be the case. Any changes made to the physical past must be consistent with the cumulative mass of observations that have been made by conscious minds. Let $O(T_m)$ be all the observed facts of the universe up to mental time T_m . Then the mental-time dependence exhibited by the physical-time series of universe states, $W(T_m)$, is subject to the constraint that $O(T_m)$ is monotonically increasing in mental time. It might be expected that this constraint would render the mental-time dependence of W inconsequential. If it allows only the unobserved physical states of the past to be changed, then — so what? The significance lies in the fact that unobserved states in the past may result in observable states in the present and the future. Furthermore, in a chaotic dynamical system (such as the brain), arbitrarily small changes to unobserved states in the past can yield macroscopically observable states in the present.

There is therefore a window of opportunity for consciousness to influence the world without violating the nomological order of the physical world. Consider some mental time T_m corresponding to physical time T_p . If consciousness at mental time T_m were to change the hitherto unobserved states of the physical world at times prior to T_p , it could bring about observable changes in the world at T_p and afterwards. Obviously, this mere logical possibility does not mean that it is realised. Nevertheless, since it is known that consciousness does influence the world, and as this is the only route by which it could do so in a nomologically bound deterministic world, then there is strong *prima facie* evidence for this theory.

3.3 Quantum-Mechanical Non-Determinism

We do not, however, live in a deterministic Newtonian world but in a non-deterministic quantum-mechanical world. In this world, the foregoing conclusion that consciousness could affect the unobserved physically-past events in order to change physically-present or future events still holds. But now it is broadened into a more general mechanism. For, in the formalism of quantum mechanics, the unobserved behaviour of a system over time is constituted by the deterministic time evolution of the Schrödinger wave, which represents the probability distribution of each observable. But the observation of a system is represented by an operator that yields a non-deterministic value of the observable, statistically weighted by the current value of the probability distribution. Thus it is not only the hitherto unobserved initial conditions that could be influenced, but each contemporary measurement of a quantum observable. There is often a confusion in popular scientific understanding between randomness and non-determinism, and the terms are often used interchangeably. Physical non-determinism refers to the fact that the successive physical states are not bound by the antecedent physical states. Randomness is the statistical property of a particular kind of non-determinism, in which a particular probability distribution is exhibited. For example, if the release of a radioactive particle from a decaying atom is non-deterministic, it means that the time of the release is not determined from antecedent physical states. If it is random, it is required to follow a certain probability distribution, namely that it has a fifty per cent probability of being released within a given length of time. Each individual occurrence is still undetermined, but collectively these events must tend toward that probabilistic pattern. (There is also confusion with a third term, 'unpredictable'. Obviously, the successive states of a physically non-deterministic system cannot be predicted from antecedent states; but, on the other hand, the successive states of a deterministic but chaotic physical system cannot be predicted, either.) I shall say that a 'purely random' physical event is one that is structured only by the physically required probability distribution. A purely random event cannot serve as a gateway to consciousness, because *ipso facto* it is not influenced by any other

constraint than the probability distribution. Nevertheless, individual non-deterministic events that are not purely random can be shaped by consciousness provided that the overall statistics tend toward the correct distribution. Suppose, as a purely hypothetical example, consciousness operated by changing the times at which neurotransmitter molecules tunneled across synapses. The timing might be adjusted to coincide with other neurochemical events occurring, to enable a particular motor pathway to be activated for some overt action to take place. Yet the statistical properties observed in that synapse may remain fully compliant with the laws of physics.

Ultimately, all initial conditions in the physical universe arose out of quantum mechanical non-determinism, for in the Big Bang the entire universe was small enough to lie within the range of Heisenberg's uncertainty. Since then, however, the universe has grown much bigger, and any observer inside the universe will not be able to observe the non-deterministic collapse of the original probability wave. Therefore there seem to be two distinct concepts of how the physical world can have an interface into consciousness: (a) through unobserved initial conditions in general, and (b) through the collapse of probabilistic wave functions. There is therefore an empirical question of whether, as a matter of fact, our consciousness uses (a) initial conditions or only (b) quantum measurements. This is still an open question but, as I shall argue elsewhere, there is tentative support for (a).

3.4 Operation of the Interface

Voluntary actions, including reports of conscious experiences, are limited to a repertoire of likely actions, surrounded (in the space of possibilities) by a penumbra of increasingly likely acts. If you ask me to visualise a colour and then report what I see, then I am likely to name a common colour, rather than a rare colour or something that is not a colour at all. So, the brain has set up a number of motor pathways, each of which leads to a particular colour utterance. And, at a common point in those pathways is a sub-network of brain cells that, depending on how they fire, could trigger any of those motor pathways. Each brain cell is generally firing at a low rate of electrical noise under

the influence of electrical noise from other parts of the brain, together with the Brownian motion of neurotransmitters in synapses and other environmental factors. Needless to say, those myriad minute influences will not normally be observed. So the necessary change is essentially a change in the collective behaviour of the brain cells due to unobserved prior conditions. On the face of it, therefore, there seems to be plenty of scope for consciousness to modify the unobserved micro-conditions so as to switch the brain's course of activity one way or another. So there does not seem to be any obvious need to invoke quantum-mechanical non-determinism.

Penrose and Hameroff have generated a lot of interest in the theory that quantum-mechanical events in the microtubules of brain cells could be the physical correlates of consciousness. There is a great deal of controversy about this, the main argument against it being that the brain is too hot for microtubules to be isolated long enough for useful coherence to develop. The line of reasoning that I have developed in the foregoing paragraphs suggests that the physical correlates of consciousness may actually be much broader.

The details of how the conscious mind interacts with the physical brain may seem far removed from our earlier broad-sweeping arguments about the nature of reality, but the two lines of investigation are intimately connected. For, if our earlier conclusion is correct, that consciousness is not reducible to physics, then we must face the challenge of giving an account of how consciousness can influence a nomologically well-behaved physical world. This is a matter for investigation over and above the study of the brain on its own.

For instance, how does the conscious mind know which non-deterministic events it must modify in order to produce its required voluntary action? If the conscious mind wants a physical outcome P at time T_P , it must trigger one or more physically non-deterministic events E_i at time $T_i < T_P$. How does it know which E_i will suffice? The most naturalistic hypothesis would be that the mind activates a neural nexus that represents the perception of the outcome, and is cross-linked to motor nerves so as to produce the necessary action – without the conscious mind needing to know such details as which muscle groups to use.

3.5 Significance of the Mind-Brain Interface

If consciousness is not derived from the physical mechanisms of the brain, then why do drugs and brain surgery affect conscious experience? A common analogical answer is to compare the brain to a television set. A magnet on the TV screen will affect the picture, but it would be absurd to infer that, say, a news presenter is inside the television set. The television set is a conduit to receive pictures from a remote studio; likewise the brain is a conduit to transceive motor commands and sensory inputs for a remote mind. A better analogue might be an email package on a computer. If your email program has been damaged in some way, then it may corrupt incoming email messages, and send out corrupt emails. Likewise a damaged brain may corrupt both imagery from the senses to the mind, and motor commands that are conveyed from the mind to the muscles, as well as access to memory and logic modules. A conceptual caveat is needed here. We talk of a correlation between the mind and 'the brain', but really there is no brain. We are really talking about a correlation between the mind and a notional construct that is abstracted from the manifold of conscious experience. Consider this trivial neurological intervention: pressing on the right-hand side of the right-hand eyeball, so as to produce phosphenes (in the left-hand visual field). According to mental monism, the finger and the eyeball are virtual. In this case, the metamind internally models the physical finger pressing the physical eyeball. From that model, it renders a visual image of the finger in the right-hand visual field along with a tactile image of the finger against the skin of the eyelid. This is the 'normal manifestation'. But the metamind will also generate a visual sensation of phosphenes in the left-hand visual field. This I shall call the 'intervention manifestation'. Thus we see what seems like a physical object causing a sensory impression, when in fact both images are brought about by the metamind.

Why are there these two different routes of manifestation? It is because consciousness has to interact with a physical world through some physical apparatus (in our case, the brain) that embeds the locus of the mental-physical interface within the physical world. So, there is inevitably the possibility that

we can interfere with the normal workings of that apparatus by sticking a finger into it, or dropping some chemical into it. Therefore the fact that there are two modes of presentation (normal and intervention manifestation) does not present a difficulty for mental monism. Why do we have brains and eyes at all? If the mind sees things by direct 'telepathic' communication with the metamind, why did eyes evolve? Almost certainly, they evolved just as evolutionary biology says they did, since the (virtual) physical world follows the laws of physics and biophysics. Nevertheless, biological structures, such as the optic nerve and the retina, are rendered in the manifest world only when they are observed directly – for example, in a dissection or invasive medical examination. They are virtual and therefore cannot be part of the mechanism that conveys information between metamental modules.

3.6 Modularity of the Metamind

We do not have access to the inner workings of the metamind, and we might never do. Nevertheless, we can attempt to reverse-engineer the metamind's computational workings by putting forward principled hypotheses and seeking to derive testable predictions from them. First, I propose that it is implausible that the metamind is a single, monolithic computation. The manifest world comprises macroscopic bodies that can be manipulated at will. If the metamind were monolithic then it would continually be subject to structural revision as the objects are built up, broken down, and rearranged. More fundamentally, it is hard to imagine how the metamind could have evolved into a single over-arching superstructure. We certainly do not want to posit some higher intelligence that designed the metamind. Therefore my first working hypothesis is that the metamind resolves into modules at the level of macroscopic objects. The existence of personal minds corroborates this hypothesis. We could imagine a universe in which streams of conscious are individuated at a smaller scale than ours are, corresponding perhaps to, say, sensory modalities. Equally, we could imagine a world in which streams of consciousness are individuated at a larger scale, corresponding to family groups, or communities, or races. That conscious minds are individuated at the level of

human beings suggests that this is a natural level of granularity in the meta-mental architecture. More importantly, however, this scale of modularity is likely to be economical. When we observe an object, we generally observe only the surface: we see the light reflected from the surface and feel the texture of the surface. So it would be otiose for the metamind to resolve into parcels within the object.

If we accept as a working hypothesis that the metamind resolves into functional modules that generally correspond to macroscopic bodies, then what can we say about the exchange of information between these modules? Consider the act of looking at a desk. The information flow in the mental monist's model is very different from that in a physical monist's model. In the physical model, the information flow is as follows. Light shines on the desk from some source such as a window, and it is reflected by the surface in all directions. Your eyeballs happen to lie in a certain position in relation to both the desk and the original light source, and a flux of light that has been reflected off the table enters your eyes. The retina converts the optical image into a two-dimensional pattern of neural signals. There is no flow of information in the opposite direction. The desk is marginally affected by the incidence of light – it will absorb a minuscule amount of heat – but it is not affected at all by your looking at it. In computing terms, the desk does not even know it is being looked at, let alone who is looking at it. Now, in the mental monist's model, the information flow is as follows. According to the hypothesis of economy, the module of the metamind that is responsible for the table will generate imagery, not for every possible direction but for only any direction in which a conscious being is observing it. Furthermore, the module needs to know about you, the perceiver. It needs to know the modelled spatial position of your eyes. And it needs to know how to address your conscious mind so that it can deliver the imagery. Conversely, your mind must at some level know how to address the desk module. So, you – the observer – send a signal of some sort to the module, requesting the imagery; and the module replies with the appropriate imagery. In crude terms, you telepathically ask the desk what it looks like, and it telepathically shows you. If mental monism is to be accepted then we do need to model the internal structure of the meta-

mind, and this seems a plausible direction in which to probe.

3.7 Concluding Remarks

I started by saying that there is at least a formal requirement to assess mental monism as a candidate solution to the mind-body problem. On examining mental monism, we find that the theory is logically supported by Berkeley's semantic argument, and it seems to be the only solution that adequately resolves the issues that Berkeley's argument raises. From a purely philosophical perspective, a strong case can be made for mental monism. But, if the metamind merely rendered a manifest world compliant with the laws of physics, then mental monism – even though correct – would be a rather sterile theory. The fact that we can report our conscious experiences demonstrates that the non-physical, phenomenal character of conscious experiences can affect the physical machinery of our nervous system. The ramifications of this simple fact are staggering, but somehow seem to slip into a lacuna of the collective intellectual awareness of the industrialised West. Lloyd (1999) has argued that, on some possible views of how the metamind operates, we can deduce testable hypotheses in the field of parapsychology. This is outside the scope of the present paper, but will be expanded in forthcoming papers.

Acknowledgements

Thanks to Dimitri Constant for his useful remarks on Wittgenstein, and Richard Conn Henry for his comments on quantum consciousness.

References

- Ayer, A.J. (1985). *Wittgenstein*. George Weidenfield & Nicolson.
- Berkeley, G. (1710). *Treatise Concerning the Principles of Human Knowledge*. Reference is to Luce, A.A., and Jessop, T.E. (Ed.) (1949). *The Works of George Berkeley, Bishop of Cloyne. Volume Two*. London: Nelson.
- Chalmers, D. (1996a). Facing Up to the Problem of Consciousness. In Hameroff, S.R., Kaszniak, A.W., Scott, A.C. (Ed.). *Toward a Science of Consciousness – The First Tucson Discussions and Debates*. Cambridge, MA: MIT Press, 5-28.
- Chalmers, D.J. (1996b). *The Conscious Mind*. Oxford: Oxford University Press.
- Cronenberg, D. (Dir.) (1999). *eXistenZ*. Toronto: Alliance Atlantis.
- Dennett, Daniel (1988). *Quining Qualia*, in A. Marcel and Bisiach, E. (Eds.). *Consciousness in Modern Science*. Oxford: Oxford University Press.
- Descartes, R. (1637). *A Discourse on Method*. Translated by Veitch, J. (1912). London: J.M. Dent.
- Eddington, A. (1928). *The Nature of the Physical World*. New York: Macmillan
- Foster, J. (1982). *The Case for Idealism*. London: Routledge & Kegan Paul.
- Goswami, Amit (1993). *The Self-Aware Universe: How Consciousness Creates the Material World*. New York: GP Putnam's Sons
- Henry, Richard Conn (2005). The Mental Universe. *Nature* 436:29
- Lockwood, M. (1987). Mind and its Place in Nature. *Lecture given at Oxford University Department for Continuing Study*.
- Lockwood, M. (1989). *Mind, Brain & the Quantum: The Compound 'I'*. Oxford: Blackwell.
- Lockwood, M. (1993). The Grain Problem. In Robinson, H. (Ed.) (1993). *Objections to*

Physicalism. Oxford: Oxford University Press.

- Lockwood, M. (1998). The Enigma of Sentience. In Hameroff, S.R., Kaszniak, A.W., Scott, A.C. (Ed.). *Toward a Science of Consciousness – The Second Tucson Discussions and Debates*. Cambridge, MA: MIT Press, 83-95.
- Lloyd, P.B. (1999a). *Consciousness and Berkeley's Metaphysics*. London: Ursa Software.
- Lloyd, P.B. (1999b). *Paranormal Phenomena and Berkeley's Metaphysics*. London: Ursa Software.
- Lloyd, P.B. (2003). Glitches in the Matrix and How to Fix Them. In Yeffeth, G. (Ed.). *Taking the Red Pill: Science, Religion, and Philosophy in the Matrix*. Dallas: BenBella Books.
- Murgenau, H. (1987). *The Miracle of Existence*. Boston: Shambhala.
- Rosenberg, G.H. (1997). *A Place for Consciousness: Probing the Deep Structure of the Natural World*. Doctoral Dissertation, Indiana University.
- Rosenberg, G.H. (1999). On the Intrinsic Nature of the Physical. In Hameroff, S.R., Kaszniak, A.W., Chalmers, D.J. (Ed.). *Toward a Science of Consciousness – The Third Tucson Discussions and Debates*. Cambridge, MA: MIT Press, 33-47.
- Rosenberg, G.H. (2004). *A Place for Consciousness: Probing the Deep Structure of the Natural World*. Oxford: Oxford University Press.
- Russell, B. (1927). *The Analysis of Matter*. London: Kegan Paul.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson.
- Smart, J.J.C. (1963). *Philosophy and Scientific Realism*. London: Routledge & Kegan Paul.
- Strawson, G. (1999). *Realistic Materialist Monism*. In Hameroff, S.R., Kaszniak, A.W., Chalmers, D.J. (Ed.). *Toward a Science of Consciousness – The Third Tucson Discussions and Debates*. Cambridge, MA: MIT Press, 25-32.

Strawson, G. (2005). *Realistic Monism: Why Physicalism Entails Panpsychism (or at least Micropsychism)*. Plenary paper, *Toward a Science of Consciousness*, 17-20th August 2005.

Von Neumann, J. (1932). *Mathematical Foundations of Quantum Mechanics*. Princeton, New Jersey: Princeton University Press.

Wachowski, L. & Wachowski, A. (Dir.) (1999). *The Matrix*, Warner Brothers.

Warnock, G.J. (1953). *Berkeley*. Harmondsworth: Penguin.

Wittgenstein, L. (1953). *Philosophical Investigations*. Oxford: Basil Blackwell. (note: Dennett references the second edition, 1958.)

Wolf, F.A. (1984). *StarWave: Mind, Consciousness and Quantum Physics*. New York: Macmillan.

Telepathy: Or, How do I Know that this Thought is Mine?

Fiona Steinkamp

1 Introduction

Telepathy is the purported ability to communicate directly mind to mind – that is, it is an ostensible exchange of information without the use of any of the known senses. Although, arguably, there is some experimental evidence for telepathy (see, e.g., Bem & Honorton 1994; Milton & Wiseman 1999; Storm & Ertel 2001), there is no evidence that it can operate either reliably or on demand. Therefore, “telepathy” here will mean “ostensible telepathy” throughout.

Conceptually, telepathy can be understood in a variety of ways. One way is to regard it as a direct mental influence on someone else’s mind. That is, one person at will *makes* another person have a particular thought. It is the sender who is the active person. For instance, someone who is lost on a mountain top may try mentally to will someone to come to where they are. Stanford (1974) has characterised this form of telepathy as MOBIA (mental or behavioural influence of an agent). Telepathy may alternatively be conceived as the reception of information about another person’s thoughts through the mediation of God or some other spiritual entity. One receives a message from God about what another person is thinking. The definition of telepathy in the glossary of the *Journal of Parapsychology* is different again. Here it is defined as “the paranormal acquisition of information about the thoughts, feelings, or activity of another conscious being”. By focusing on the *acquisition* of information, this definition implies that the receiver is the active person – for instance, someone waiting to hear about the outcome of an