Fundamental Consciousness: Panpsychism, Priority Cosmopsychism and Information

By

Nino Kadić

Submitted to

Central European University

Department of Philosophy

In partial fulfilment of the requirements for the degree Master of Philosophy

Supervisor: Philip A. Goff

Budapest, Hungary

2016

Abstract

In this thesis, I argue that consciousness cannot be explained solely in terms of the facts revealed to us by empirical science. This is explained through a discussion of the explanatory gap between our knowledge of physical facts and those regarding consciousness. Specifically, I argue against the Phenomenal Concept Strategy, a fairly recent approach to answering anti-physicalist arguments. I then claim that if we take consciousness as something that cannot be explained in purely physical terms, we must also say that it is a fundamental and ubiquitous feature of reality. I explain the motivation for accepting this claim by considering what it means for something to be an intrinsic nature and by addressing the problem of strong emergence. The result of these considerations is panpsychism, the theory that consciousness truly is fundamental and ubiquitous. After that, I introduce the combination problem, the issue of how small subjects combine into bigger subjects. Priority cosmopsychism is then offered as a possible way of avoiding the combination problem. I then argue against priority cosmopsychism, claiming instead that we should remain panpsychists. I then move on to suggest a solution to the combination problem based on a discussion of syntactic and semantic information. I argue that information can be used to explain how small subjects bond to form bigger subjects. The final result is an informational form of panpsychism.

Acknowledgments

I would like to thank my supervisor Philip A. Goff for offering invaluable comments and for guiding me through the whole process of writing and submitting this thesis. I am also grateful to the second reader Michael V. Griffin for his patience and help. Olesya Bondarenko, Garrett Mindt and Anthony Skoviak, my friends and peers at Central European University, have all provided helpful comments and tips during the process of writing. I extend my gratitude to them as well. I have discussed parts of this thesis with Yujin Nagasawa and Khai Wager from the University of Birmingham, so I wish to thank them for their insightful feedback. Most importantly, I would like to thank my parents, Katarina Butković and Šerif Kadić, for always supporting me and being there for me. I dedicate this thesis to them.

Table of Contents

Introduction	1
1. Consciousness, the Epistemic Gap and Phenomenal Concepts	3
1.1 The Epistemic Gap and Physicalism	3
1.2 The Extreme Epistemic Gap and Subjectivity	7
1.3 The Phenomenal Concept Strategy	11
2. Panpsychism, Intrinsic Nature and Emergence	16
2.1 Panpsychism and Intrinsic Nature	16
2.2 Initial Objections to Panpsychism	18
2.3 The Problem of Radical Emergence	20
3. The Combination Problem and Priority Cosmopsychism	25
3.1 Panpsychism and the Combination Problem	25
3.2 Priority Cosmopsychism and the Derivation Problem	27
3.3 Objections to Priority Cosmopsychism	30
4. Phenomenal Bonding, Information and Panpsychism	35
4.1 The Phenomenal Bonding Solution	36
4.2 Chalmers' Dual-Aspect Theory of Information	39
4.3 Seager on Semantic Information and Consciousness	41
4.4 A New Dual-Aspect Theory of Information	43
4.5 Consciousness and Grounding	45
Concluding Remarks	48
References	49

Introduction

Can consciousness be explained in purely physical terms? This is one of the most prominent questions in current philosophy of mind. On the one side, there are anti-physicalists, who find it hard to believe that our immensely rich and complex mental life can be explained completely by appealing to the facts revealed to us by physics. On the other side, there are physicalists, who think that it must be physically explainable because it would be equally hard to believe that everything *except* consciousness can fit a physicalist framework. Then, there are panpsychists like me, who claim not only that consciousness cannot be reduced to the physical, but that it is also a fundamental and ubiquitous feature of the universe.

To defend this claim, I will largely focus on recent philosophical literature. There has been a revival of interest regarding accounts of consciousness that see it as fundamental and ubiquitous, such as panpsychism. One of the main motivations for such accounts is the idea that matter must have an intrinsic nature, as discussed by philosophers such as Bertrand Russell, Alfred North Whitehead and, more recently, David Chalmers, Galen Strawson, Philip Goff and others. The claim I make in this thesis is that accepting that consciousness cannot be reduced to the purely physical leads us to panpsychism. However, panpsychism is faced with the notorious combination problem: how do small subjects combine to form larger subjects? The contribution I offer in this thesis is a new approach to tackling this difficulty by relying on a discussion of information. I argue that we can explain how subjects combine in terms of information and thus avoid the combination problem. In Chapter 1, I argue that an explanatory gap exists between the knowledge we have of physical facts and the knowledge we have regarding consciousness. Specifically, I focus on a fairly recent physicalist approach that attempts to dispel this gap by arguing that it only exists because of the differing concepts that we use to describe physical and phenomenal facts. Based on the work of David Chalmers, I argue that this approach fails and that we should reject physicalism.

In Chapter 2, I claim that accepting that consciousness is irreducible to the purely physical leads us to panpsychism, the theory that consciousness is fundamental and ubiquitous. Focusing on research done by Galen Strawson and Thomas Nagel, I argue that radical emergence, where a new property emerges from something that does not originally instantiate it, is highly improbable. From this, I argue that consciousness is weakly emergent, meaning that it is already present at the lower level from which it emerges.

In Chapter 3, I introduce the combination problem, a serious objection to panpsychism. This problem raises the question of how smaller subjects combine to yield larger subjects. I focus on priority cosmopsychism, a recent approach discussed by Yujin Nagasawa and Khai Wager that claims to avoid the combination problem. I argue that panpsychism is preferable because priority cosmopsychism is committed to deeply problematic ontological claims.

In Chapter 4, I focus on the phenomenal bonding solution to the combination problem offered by Philip Goff, who argues that small subjects can be in certain relations that bring about larger subjects. With reference to David Chalmers' and William Seager's work on information, I propose a way of making sense of Goff's solution by arguing that semantic information can explain how less complex subjects combine into more complex subjects.

1. Consciousness, the Epistemic Gap and Phenomenal Concepts

At the heart of debate in the philosophy of mind is the question of whether or not phenomenal consciousness is something over-and-above the purely physical. I will argue that it is and that it cannot be reduced solely to the physical facts of the universe. Phenomenal consciousness is experience that involves a qualitative aspect for the subject having that experience. In other words, there is something it is like for the subject to be in that state. There seems to be a gap between our knowledge of this aspect of experience and our knowledge of physical facts. The result of this, as will be argued, is that physicalist theories of consciousness cannot provide an adequate account of phenomenal consciousness. However, fairly recently, the Phenomenal Concept Strategy was proposed in order to provide an answer to anti-physicalist arguments. While this new strategy certainly is promising, I will demonstrate that it fails to account for the qualitative aspect of conscious experience by relying on the Master Argument introduced by David Chalmers. After addressing an important objection to his argument raised by Esa Díaz-León, I will conclude that her objection misses the mark and that we have good reasons to think that phenomenal consciousness really cannot be accounted for in terms of the purely physical.

1.1 The Epistemic Gap and Physicalism

I understand physicalism in the strict sense as the view that every phenomenon of reality can be explained in terms of the facts revealed to us by physics. The strict physicalist argues that there is nothing over and above these facts. Galen Strawson (2006) called this view 'physicSalism' to distinguish it from "real physicalism", the position that consciousness should still be understood as physical, but that more is needed than the facts of physics to fully capture the nature or essence of conscious experience. I will argue that phenomenal consciousness cannot be explained in purely physical terms in the strict sense mentioned above. It is important to emphasise that I see physicalism as having the aim of being an *explanatory* framework, so if a physicalist account of consciousness fails to explain it in terms of such facts, it is to be rejected. This will be important later on in this chapter.

As I will argue, physicalist accounts indeed fail to explain consciousness purely through the facts of physics. The potential reason of their failure is the observation that there is an epistemic gap between our knowledge of physical facts and our knowledge of phenomenal facts. One way of demonstrating the existence of the epistemic gap was discussed by Frank Jackson (1982), who argued that there is a form of knowledge in addition to knowing all the physical facts about reality. This is the now-famous knowledge argument. Jackson prompts us to imagine a scientist called Mary who spent her whole life in a black-and-white room researching all the facts regarding colour (1982: 130). She is the perfect colour scientist, meaning that she knows all the possible physical facts about colour that there is to know. However, there is something missing here: her knowledge of what it is like to experience, for instance, the colour red. Mary has never seen the colour red, so when she exits the room and sees red for the first time, she will gain additional knowledge. She will acquire the knowledge of what it is like to perceive the colour red. This knowledge cannot be captured, as the argument demonstrates, by the full index of physical facts about colour. In other words, Mary's knowledge of all physical facts does not entail the knowledge of what it is like to be experiencing colours. The conclusion of the argument is that not all information is physical information (Nida-Rümelin 2009). There is something beyond the complete set of physical data, namely the knowledge of qualitative experience, that physicalism fails to account for. Thus, physicalism is false, as it cannot explain this aspect of consciousness in purely physical terms (Jackson 1982: 130).

Thomas Nagel (1974) cast this problem in different terms. He argued that we can never know what it is like to be a creature with a different conscious experience. He used the bat as a particularly strong example, considering how dissimilar our perceptual apparatus is from the one of bats. We can dissect the bat, study how it uses echolocation, document its behaviour and make predictions. We could even map out the bat's neurosystem and predict what the bat experiences or how it will behave. However, we could still not know what it is like to be a bat, having those experiences from its own point of view. One could imagine an even more alien creature, whose perceptual apparatus is completely different from ours, but the same logic applies to humans as well. Every experience we have passes through the 'filter' of subjectivity. An experience is had by a subject in a unique manner. The only way for a subject to experience what another subject experiences is by becoming that subject, but then the original subject is lost, so there is ultimately only one subject left. This is, of course, an absurdity. Hence, there is a unique and qualitative "what's-it-like" aspect of being a conscious subject, elusive to purely physical descriptions. What Nagel's argument demonstrates is that no amount of external, third-person knowledge can help us understand what it is like to be the bat, and what it experiences from its internal subjective point of view (Van Gulick 2014). If this is true, subjective experience cannot be explained in terms of the facts of physics. This is a reason to reject physicalism.

David Chalmers has also argued for the existence of an epistemic gap between the complete knowledge of every physical fact about the universe and the phenomenal experience of conscious beings. He takes P to be the complete microphysical truth about the universe, a long conjunction of every microphysical property of every fundamental microphysical entity, and Q to

denote an arbitrary truth about phenomenal consciousness (e.g. that there is somebody who is phenomenally conscious) (Chalmers 2007: 168). The problem here is that one can know P but fail to know Q, like in the case of Mary. In other words, knowledge of Q is not deducible by a priori reasoning from knowledge of P, so there is an epistemic gap between the knowledge of the complete set of physical facts and knowledge of phenomenal experience (Chalmers 2007: 168– 9).

This becomes clearer when we examine Chalmers' treatment of the conceivability of zombies. The notion of conceivability he uses is that of *negative ideal conceivability* (Goff forthcoming (a)). A sentence S is *negatively conceivable* if it is not a priori true that S is false (Goff forthcoming (a)). A sentence S is *ideally conceivable* when it is conceivable under the assumption of ideal rational reflection (Chalmers 2003: 119), that is, in conditions when the contingent limits of human reason are not taken into account (Goff forthcoming (a)). A sentence S is *negatively ideally conceivable* when "*no amount of a priori reasoning* could rule out its truth" (Goff forthcoming (a); emphasis in the original).

In Chalmers' argument, zombies are not sluggish undead cannibals, but creatures that are our exact atom-for-atom duplicates, with the exception that they lack any sort of phenomenal consciousness. A more extreme form of this argument would be to conceive of a zombie world which is completely identical to ours in terms of physical facts, but where phenomenal consciousness does not exist at all. Chalmers argues that the conceivability of zombies entails that they are metaphysically possible, which results in the conclusion that consciousness is nonphysical (2003: 106). He puts the argument forth in a more formal manner as follows (Chalmers 2003: 106): (1) It is conceivable that $P \land \neg Q$.

(2) If it is conceivable that $P \land \neg Q$, it is metaphysically possible that $P \land \neg Q$.

(3) If it is metaphysically possible that $P \land \neg Q$, then materialism is false.

(4) Materialism is false.

Put simply, if we take two worlds where one world has both P and Q, while the other has P, but not Q, then it seems that P does not metaphysically necessitate Q. If it is accepted that conceivability entails metaphysical possibility, which is admittedly controversial, then Chalmers' argument is successful in demonstrating the existence of an epistemic gap since something more is needed to explain consciousness, over-and-above P or the complete set of physical facts¹. If this set of physical facts cannot account for consciousness, physicalism (or materialism) is false.

1.2 The Extreme Epistemic Gap and Subjectivity

I think that the arguments presented above are successful. The qualitative aspect of subjective conscious experience truly cannot be explained in physicalist terms. If subjectivity involves privileged access, meaning that only the experiencing subject has direct access to their own conscious states, the physicalist programme's aim to explain consciousness in terms of physical facts renders the concept of subjectivity obsolete. Negating that subjects have phenomenal "what's-it-like" consciousness results in denying that any form of subjectivity exists in the world, considering that all the data about the universe will then be explicated in third-

¹ Chalmers argues that an ontological gap can be inferred from the epistemic gap: "if we cannot explain consciousness in terms of physical processes, then consciousness cannot be a physical process" (Chalmers 2007: 169). The conceivability of a consciousness-less but otherwise physically identical world to our own means that the existence of consciousness is a further ontological fact of reality not entailed by the complete set of physical facts (Chalmers 2007: 169). However, in this paper, I will focus exclusively on the epistemic gap.

person, objective terms. What we commonly take to be internal subjective experience should thus be understood, as the physicalist argues, as another objective third-person fact or set of facts about the world. But how could this be the case? This seems, at least intuitively, as a point that could be raised against physicalism because it involves the strange claim that subjective facts do not exist in reality. Surely, it is conceivable that a subject might be having a private thought about a memory that makes them feel nostalgic and melancholic. Assuming that the subject does not express this thought in any way, entertaining it for a short time and then forgetting about it, I fail to see a way how we could consider this particular experience as something that can be captured in physical terms. The physicalist might argue that, perhaps in the future, we could scan the brain of the mentioned subject and see a simulation of the experience on a supercomputer, or even transfer it and "play" it in the brain of another subject who would then have the same subjective experience. There would then be a sense of how we could understand internal subjective experience as explicable in objective terms.

Putting aside the question of whether this technology is possible, I want to deny that the simulation, whether on the supercomputer (external) or within the brain of a different subject (internal), gives us any relevant, let alone complete epistemic knowledge of the original subject's conscious state. In the first case, where the simulation is external (e.g. projected on the screen of the supercomputer), there is no subject that is having the original experience. The people who are watching the original subject's experience on the screen are not experiencing a phenomenal state, but, quite crudely, a simulation of images and sounds. While watching the supercomputer screen, they are observing what is external reality to them, but what was internal reality to the original subject. This is different from the case of the internal simulation both ontologically and epistemologically, as both the way of perceiving the experience and the way the experience is

realised are different. This option is a non-starter for physicalists who want to argue that we can know what it is like to be another subject and that conscious states can be captured in thirdperson objective terms.

The second case, of the internal simulation, is much more interesting. Here the original subject's (S1) experience of nostalgia and melancholy is simulated or directly transferred to the brain of another subject (S2). There are several ways to argue that the simulation does not provide S2 with insight on what it is like to be S1. To eliminate worries of time and space, let us assume that S2 is experiencing the simulation at exactly the same time as S1 and that the simulation correctly reproduces the experience of S1. While this simulation certainly gives S2 a very personal look into the mental life of S1, it does not capture what it is like to be S1 simply because S2 is another subject. Returning back to the 'filter' of subjectivity I mentioned earlier, it is clear that the experience of S2 will have an additional layer the original experience does not have, namely the qualitative aspect of being S2. Hence, S2 will still experience the simulation, no matter how intricate and advanced it is, as *their own* experience. A useful analogy might be reading a poem or a novel. Even if the reader had access to the mental life of the author, the final experience of the work will still be the reader's. If the "what's-it-like" aspect of being S1 is simulated within S2, it is impossible for the simulation to be completely accurate, as "being S1" is something that can only be accurately experienced by S1. As I have argued earlier, a complete simulation would require S2 to become S1, but this is nonsensical. Thus, there is no way one could account for subjective conscious experience in objective terms, even when we consider hypothetical examples of future technology.

The point of this exposition is to illustrate that the epistemic gap between the knowledge of physical and phenomenal facts is *extreme*, a gap not to be ignored by any serious account of

consciousness. I wanted to show that even in the case of the direct perception of a subject's experiential state, which is plausibly the most intimate way we can get to know someone, the knowledge gained still does not help us to fully understand what it is like to be that subject. Franz Brentano made a similar point, worth being quoted at length:

"In fact, it would be like saying that, of course, neither a blind man nor a deaf man could compare colors with sounds, but if one sees and the other hears, the two together can recognize the relationship. [...] Obviously, it makes no difference whether the blind man and the deaf man are far apart or near one another. If they lived together permanently in the same house, indeed, even if they had grown up together as inseparably as Siamese twins or even more so it would not increase the possibility of the hypothesis one bit. Only if sound and color are presented jointly, in one and the same reality, is it conceivable that they can be compared with one another." (1973: 123)

Brentano uses this passage to show that phenomenal experiences cannot combine, but I think it can also be interpreted to show that, no matter the circumstances, the blind man cannot have the deaf man's knowledge of the experience of seeing, and vice-versa. When the experience of S1 is projected onto a computer screen or even directly into the brain of S2, it does not become a phenomenal fact, but remains a third-person, external physical fact of reality. The only phenomenal facts in this scenario are the isolated subjective qualitative experiences of S1 and S2. The second subject can never have knowledge of what it is like to be the first subject, even if the knowledge in question is of an extremely intimate and 'high quality'.

1.3 The Phenomenal Concept Strategy

As a possible response to the problems for physicalism resulting from the epistemic gap, I wish to focus exclusively on the phenomenal concept strategy (PCS), a fairly recent and promising physicalist response to anti-physicalist arguments. Proponents of PCS accept that an *epistemic* gap exists, but do not accept that an *ontological* gap follows from it. Moreover, they do not see the epistemic gap as a gap of *knowledge* between physical and phenomenal facts. Rather, the gap exists because of the difference in the *concepts* we use for consciousness and those we use for physical processes (Chalmers 2007: 167). There is no epistemic gap, however, between consciousness and physical processes in reality. There is only a disconnection between how we describe conscious experiences and features of the physical world (Chalmers 2007: 167). That is all there is to the explanatory gap^2 . Phenomenal concepts are those we employ when we think about the qualitative subjective aspect or "what's-it-likeness" of our conscious experience. For example, the phenomenal concept "red sensation" denotes a specific sensation that occurs when we see red things (Stoljar 2005: 469). As David Papineau puts it, physicalists take phenomenal concepts to be distinctive "only at a *conceptual* level", referring to nothing but material properties "that can also be referred to using nonphenomenal material concepts" (Papineau 2007: 111). So the proponent of PCS blocks the move from an epistemic to an ontological gap and defines the epistemic gap as holding between concepts, rather than pertaining to our knowledge.

But can the proponent of PCS get away with this strategy? I believe it is now time to consider the master argument, an objection to PCS introduced by David Chalmers (2007). According to Chalmers, the PCS proponent puts forth a thesis C to attribute certain (key) psychological features to humans, arguing that if C is true: a) humans really do have these

² I use the terms 'explanatory' and 'epistemic' gap interchangeably.

features, b) C explains the epistemic situation regarding consciousness, and that c) C is explicable in physical terms (Chalmers 2007: 172). The goal of the PCS proponent is to provide a physicalist explanation of how the explanatory gap arises, while denying the existence of an ontological gap. Chalmers aims to show either that PCS is not physically explicable, or that it fails to explain our epistemic situation regarding consciousness (2007: 173). His argument goes as follows. The complete index of physical facts about the universe is P. The aforementioned PCS thesis is C. Chalmers argues either that 1) if P&~C is conceivable, C cannot be explained in physical terms, or that 2) if P&~C is inconceivable, then C fails to explain our epistemic situation concerning consciousness (Chalmers 2007: 174). Thus the PCS argument is faced with a destructive dilemma.

The first horn of the dilemma is similar to the standard zombie conceivability argument mentioned earlier. If one can conceive of our atom-for-atom duplicates that lack the psychological features of C, then a real epistemic gap between the complete physical facts P and the C thesis will exist: "there will be no wholly physical explanation that makes transparent why thesis C is true" (Chalmers 2007: 174). So if we can imagine a world identical to our own, but with no psychological features of C, while knowing that we have those features in this world and the same physical structure that our imagined counterparts have, it seems that something additional is necessary to account for thesis C. Since proponents of PCS typically accept that this is conceivable, the first horn presents a serious problem to them (Chalmers 2007: 175).

The second horn of the dilemma is a bit more complex. It presents a problem for the PCS proponent because it forces them to concede that C cannot account for our epistemic situation with regard to consciousness. It goes as follows: if it is inconceivable that there is a world physically identical to ours, but without the psychological features of C, then even zombies can

satisfy C (Chalmers 2007: 176). However, if zombies do not share our epistemic situation but satisfy C, then C cannot explain our own epistemic situation (Chalmers 2007: 176). If P&~C is inconceivable, the C must necessarily follow from P. If that is true, then zombies can satisfy C as well, as P is all that is necessary for the psychological attributes of C to be instantiated. But they have, by definition, a different epistemic situation, so C fails to explain the difference between the epistemic situation of zombies and humans. Again, since proponents of PCS typically grant the conceivability of zombies, the second horn of the dilemma presents a problem for PCS as well.

Recall the definition of physicalism introduced at the beginning of this chapter. Physicalist theories are committed to explaining our epistemic situation regarding consciousness in terms of the purely physical. If PCS cannot do that, as Chalmers argued, then it should be rejected. However, Esa Díaz-León (2010) has argued that Chalmers misconstrues PCS in his master argument. She focuses on the second horn of the dilemma and argues that it does not pose a serious challenge to PCS even if we accept that zombies do not share our epistemic situation (Díaz-León 2010: 943). According to her, the explanatory goal of PCS is not to explain our entire epistemic situation with respect to consciousness, but only to offer an alternative explanation of the epistemic gap in terms of concepts, thus blocking the potential move to an ontological gap between the physical and the phenomenal (Díaz-León 2010: 943–4). Díaz-León argues that we could define the psychological features ascribed by C in purely functional terms, so that a zombie could also have concepts that play different roles to them (2010: 945). Since concepts used by zombies also play such differing roles, sentences involving these concepts will not be a priori inferable from sentences involving only physical concepts (Díaz-León 2010: 945). In other words, Díaz-León denies that C must necessarily follow from P, as stated in Chalmers' second horn. If that is true, zombies do not present a problem for the claim that C can explain the epistemic gap (Díaz-León 2010: 945).

Díaz-León's suggestion strikes me as unsatisfying. In defining the aim of PCS as explaining the epistemic gap in terms of an inferential disconnection, rather than our epistemic situation regarding consciousness, she does not address the reasons for why there is an epistemic gap in the first place. As Chalmers argued, even though PCS might explain the inferential disconnection between physical and phenomenal *beliefs*, the crucial point of the anti-physicalist is the epistemic gap between physical and phenomenal *knowledge* (2007: 183–4), so that PCS must explain why we can conceive of beings that are physically identical to us, but lack phenomenal states of which *we have* substantive knowledge (Díaz-León 2010: 946). It seems that Díaz-León gives up on the aim of explaining our epistemic situation with respect to consciousness, effectively reducing PCS to merely offering a different definition of the epistemic gap (and stopping there)³.

After introducing the well-known arguments used to establish the existence of the epistemic gap, I have argued that said gap is indeed quite extreme. The phenomenal concept strategy was then introduced as a possible way of explaining the epistemic gap in different terms, so as to deny that it is a gap in knowledge of physical and phenomenal facts and thus to stop the move from an epistemic to an ontological gap (as some philosophers try to do). However, this strategy was faced with the destructive dilemma, raised by David Chalmers, of either not being explicable in purely physical terms or of failing to explain our epistemic situation regarding

³ A more extreme claim would be that accepting this aim of PCS, to merely account for different ways we describe a monadic physical reality, leads us to the view that the world is radically unintelligible, as we would have no way of accounting for how numerous ways of carving up the world 'hang together' (Goff forthcoming (a): 102). In other words, the necessary connections between physical and phenomenal concepts would be unintelligible in that case (Goff forthcoming (a): 102).

consciousness. The objection posed by Esa Díaz-León to Chalmers' argument was that it misconstrues the aim of PCS. As I have hopefully shown, this objection is unsatisfying as it reduces the PCS strategy to a weaker claim. Thus, for now, I believe that we have a good reason to believe that phenomenal consciousness truly is inexplicable in purely physical terms.

2. Panpsychism, Intrinsic Nature and Emergence

After concluding that consciousness cannot be reduced to the purely physical, I want to argue for panpsychism, the theory that consciousness is a fundamental and ubiquitous feature of reality. While this jump may seem as a bit extreme, it is surprisingly easy considering the premises involved. I will first introduce panpsychism and explain why I choose to focus on it rather than on other theories such as substance dualism and panprotopsychism. Then, one of the main reasons for taking panpsychism seriously will be explained: the notion that there must be an intrinsic nature to matter. After that, some initial objections to panpsychism will be addressed, such as its apparent counterintuitivity and lack of empirical reasons for believing in it. Finally, I will argue against radical emergence, a form of emergence where a new property comes about from something that does not already possess that property. I will show radical emergence to be highly improbable, if not impossible. It will be demonstrated that if we take consciousness to be irreducible to the purely physical, we must also commit to the claim that it is fundamental and ubiquitous. In other words, the result of these considerations will be a defence of panpsychism.

2.1 Panpsychism and Intrinsic Nature

Panpsychism is generally understood as the theory that consciousness is a fundamental and ubiquitous aspect existing throughout the universe (Seager 2010). The theory has a rich history in philosophy. It was entertained by early philosophers such as Thales, who developed an argument for the attribution of minds resembling panpsychism, Anaxagoras, who is interpreted sometimes as proposing that everything has some portion of mind, and Heraclitus, whose *logos*, the principle of the universe, could be interpreted as reminiscent of panpsychism (Seager 2010). It is also an important part of certain spiritual and religious traditions, such as in Taoism where the *Tao* (usually translated as 'way' or 'path') is postulated as the cosmic principle, or in the Buddhist schools Huayan and Tiantai, where the Buddha nature is attributed to both living and non-living beings. It persisted throughout the centuries, in some form or another, occupying the minds of the likes of Giordano Bruno, Baruch Spinoza, Georg W. F. Hegel and so on. Nowadays, it took a more modern form informed by recent discussions in the philosophy of mind. Some current philosophers who have discussed panpsychism are Thomas Nagel, David Chalmers, William Seager, Galen Strawson and Philip Goff.

It is important to specify why I exclusively focus on panpsychism, as there are other theories which take consciousness to be fundamental, ubiquitous or both, in one way or another. The reason why I will not take into account substance dualism, the view that there is a distinct physical or material substance and a distinct phenomenal or mental substance, is that I wish to focus on monistic theories of consciousness, that is, those which do not posit the existence of more than one substance. Substance dualism is also an unpopular view nowadays, so it will not be discussed further. The other option I will not take into account is panprotopsychism, which takes proto-phenomenal properties to be ubiquitous and present at the fundamental level. Protophenomenal properties are those which are not phenomenal, but that can "collectively constitute phenomenal properties, perhaps when arranged in the right structure" (Chalmers 2015: 259). There are two reasons why I will not discuss this view. The first one is that panprotopsychism introduces a further problem of accounting for what these proto-phenomenal properties are, as well as the problem of explaining how proto-phenomenal properties can constitute phenomenal properties.

The second reason is based on considerations of intrinsic nature. Indeed, one of the most important reasons for accepting panpsychism in recent literature rests on ideas introduced by Bertrand Russell, who argued that empirical observation only provides us with the mathematical structure of matter, or, stated differently, the spatiotemporal and behavioural facts about matter, without actually saying anything about what matter really is intrinsically (Russell 1927/1992: 10, 384). In other words, the information we gain on physical properties is incomplete as it does not reveal the inner nature of matter. For example, take the physical properties of mass, spin, electric and colour charge. These properties are defined by physics only in terms of how they relate to other properties. Physics is silent about what these properties are intrinsically, or what they are grounded on. Alfred North Whitehead had a similar argument when he emphasised that there must be a deeper underlying nature behind matter in addition to how it appears to us (1967: 133). Since the only intrinsic nature we know of is consciousness, it is more parsimonious to postulate it rather than intrinsic natures of a different kind. As William Seager argued, the only justification for postulating proto-phenomenal properties seems to be the verbal demand that they are 'non-mental', which is unjustified considering that we already have one intrinsic property "ready to hand" for that (2006: 137). Therefore, I will not consider panprotopsychism in this paper.

2.2 Initial Objections to Panpsychism

Despite its long history and the big names behind it, panpsychism still strikes most philosophers as a deeply counterintuitive idea. The thought that, say, an atom could be conscious sounds to many like an absurdity. The panpsychist view of the world, where everything is essentially conscious in one way or another, does not correspond to our general intuitions. This is one of the first reactions and objections people have to panpsychism, but I think that it is unjustified. Intuitions, while useful, cannot be taken as leverage points to discard a whole theory. Furthermore, it is not true that people *generally* share the intuition that panpsychist-like theories are implausible. In some major Eastern philosophical and spiritual traditions, such as Taoism and Shintoism, the idea that consciousness permeates everything is a fairly common belief. Similar beliefs can be found in certain groups of indigenous people. Moreover, throughout history, intuitions have changed. Therefore, I think that the intuitions one might have about panpsychism should not be used as a motivation to criticise or abandon the theory. Perhaps panpsychism could motivate a re-evaluation of our general intuitions instead. It might turn out to be a more intuitively plausible theory once the issue of radical emergence is explained, which will be done later on in this chapter.

One further worry is that there are no empirical reasons for taking panpsychism seriously. When we observe non-living objects around us, there is no justification for ascribing any form of consciousness to them. But this worry goes even further: if consciousness is a fundamental aspect of reality, then why do we never see matter behaving in a conscious-like way on the fundamental level? The lack of evidence for panpsychism is thus apparent on two levels, on the one we perceive in daily life and on the one we observe through empirical science. There are two ways to provide an answer to this objection. In the first case, regarding our day-to-day interaction with objects, the panpsychist is not really faced with a pressing issue, as they admit that consciousness comes in degrees. So the fact that a rock does not display conscious behaviour might simply indicate that a sufficient structure for the emergence of a more complex, larger subject or macrosubject is not met in its case, meaning that the rock is a mere aggregate of the smallest physical entities or microsubjects that compose it. Explained in simple terms, a panpsychist is not committed to the claim that *everything* is conscious. A panpsychist might claim that only the smallest entities that compose everything are conscious and that only some configurations of those entities will result in a more complex subject. It might be that only humans, some higher animals and physical ultimates are truly conscious, while other entities exist as mere aggregates of basic parts.

In the second case, where there is no evidence for conscious-like behaviour on the fundamental level, the answer might be that, if consciousness is fundamental, the efficacy it exerts within the processes of the universe would be undetectable. To us, its efficacy would just be a part of the laws of nature. It might be that our concepts regarding the laws of nature and causality are wrong in describing the universe as mechanistic rather than consciousnessinvolving. Perhaps the laws we find in nature truly are manifestations of consciousness, but we have failed to identify that. If the question is why these processes do not behave freely or even erratically as self-conscious subjects do, then the answer could be that the primitive consciousness of physical ultimates is not complex enough to behave in such a way. But why postulate that consciousness necessarily involves degrees of freedom? It might be that consciousness on every level is determined, with free will being only an illusion. Humans, as beings who think of themselves as having free will, could have made a mistake in identifying themselves as free and the laws of nature as determined. It might be that either determinism or free will is the case on both levels. The strategy of accepting determinism on every level would at least be in line with the general preference for forms of determinism prevalent today.

2.3 The Problem of Radical Emergence

A positive reason for why we should take panpsychism seriously will now be explained through a discussion of weak and strong (or radical)⁴ emergence. Weak emergence is a form of emergence where the higher-level properties of a system can be wholly explained in terms of the structure and activity of lower-level processes (Vintiadis (no date)). A commonly used example

⁴ The terms "radical" and "strong" are not necessarily interchangeable, but in the context of this discussion it is unproblematic to treat them as synonyms.

of weak emergence is liquidity. It initially appears that the liquidity of water, for instance, is a feature that is not present at the level of the molecules that compose it (Strawson 2006: 13–4). However, we can easily understand that this is not the case, as liquidity simply *is* a certain molecular structure, where molecules are not bound together, but slide past each other in accordance to van de Waals' laws of molecular interaction (Strawson 2006: 13). So, in the case of weak emergence, higher-level properties (such as liquidity) are in principle ontologically and causally reducible to the lower level (such as molecular structure), meaning that they can be completely explained in terms of the properties and relations found at the lower level (Vintiadis (no date)). A successful weak emergentist theory of consciousness would be capable of explaining it as a phenomenon solely in virtue of the features present at the lowest, or fundamental, level of reality.

In the case of strong (or radical) emergence, however, the higher level is completely ontologically distinct from the lower level in the sense that it *cannot* be explained in terms of the facts of the lower level. Regarding consciousness, as Galen Strawson puts it: "The idea will be that X remains *in itself* wholly and utterly non-experiential, but *gives rise to* something wholly ontologically distinct from itself, i.e. Y" (Strawson 2006: 23). In this case, there is an *ontological novelty* present at the higher level, i.e. a property which is not present at the lower level and cannot be explained in virtue of the properties present at the lower level. Most importantly, these two types of emergence should be understood in terms of intelligibility: a) in the case of weak emergence, the higher-level properties are in principle deducible from the lower-level domain, b) in the case of radical (or strong) emergence, higher-level properties are in principle *not* deducible from the lower-level domain (Chalmers 2006a: 244). In this context, 'deducible' means that it is possible to calculate or 'figure out' the properties at the higher level from the knowledge we have

of the properties at the lower level. It should be noted that there are no real cases of strong (or radical) emergence that have ever been observed, which is a point conceded even by those who defend its existence (Vintiadis (no date)).

If we assume that phenomenal consciousness cannot be reduced to the purely physical, then it must be a fundamental feature of the universe. In other words, if we grant that consciousness emerges from lower-level organisations of matter, then it can only emerge in the weak sense of the term. Strong or radical emergence, where x emerges from y as an ontological novelty without previously being entailed by y, is an *ex nihilo* relation where 'something comes out of nothing'. It is hard to see how this could happen. Indeed, in addition to the earlier point that there are no observed cases of such emergence, some have directly argued that it is impossible. For example, Strawson describes it as 'magic' and as something "not even God can understand" (2006: 24). Thus, if consciousness is irreducible, we must assume that it is weakly emergent. Strawson explained this guite well: "For any feature Y of anything that is correctly considered to be emergent from X, there must be something about X and X alone in virtue of which Y emerges, and which is sufficient for Y" (Strawson 2006: 18). Therefore, if consciousness (Y) cannot be identified with the purely physical (X), then it can be an emergent feature of X in two ways: radical emergence or weak emergence. Since we reject radical emergence as an impossibility, it can only be emergent in the weak sense. This means that consciousness must already be present on the fundamental level of reality.

Thomas Nagel made a point similar to Strawson's: "There are no truly emergent properties of complex systems. All properties of a complex system that are not relations between it and something else derive from the properties of its constituents and their effects on each other when so combined" (Nagel 1979: 182). This is one of the premises he uses to argue for

panpsychism. The other premises are: a) there are no immaterial souls, b) living organisms, including humans, are complex material systems, and c) mental states are not physical properties, but they are properties of an organism (Nagel 1979: 181–2). If consciousness is not implied by the physical properties of an organism, but must come about from the constituents of the organism, then these constituents must have *nonphysical properties from which consciousness emerges under certain conditions* (Nagel 1979: 182). Nagel then postulates that all matter can form an organism or be a part of a conscious system because it can be broken down to a fundamental enough level (perhaps that of quarks) and integrated into consciousness-involving systems (1979: 182). In other words, all matter must have these nonphysical properties. Thus, consciousness is not only a fundamental, but also a ubiquitous feature of reality, which is the theory of panpsychism.

If one accepts that consciousness is fundamental, but denies that it is ubiquitous, then one ends up with micropsychism, the position that only *some* physical ultimates are consciousnessinvolving while others are not (Strawson 2006: 24–5). In addition to Nagel's claim that all matter can be a part of a conscious system, a further reason for why we should reject micropsychism and accept panpsychism was offered by Strawson. He argues that the claim that only some physical ultimates are conscious while others are not is like saying that only some but not all physical ultimates are spatiotemporal (Strawson 2006: 25). Both Nagel and Strawson's reasons for preferring panpsychism over micropsychism strike me as highly plausible. Hence, there are good reasons to think that if there are some consciousness-involving physical ultimates, then all physical ultimates must be consciousness-involving. This is strengthened by the generally wellreceived notion that mental properties can be realised by a wide variety of material realisers. For example, I feel pain when someone punches me, but a duplicate of me whose cognitive apparatus is significantly different from mine, such as a squid's neurosystem or an alien one based on silicon, could presumably experience pain as well. Considering all that has been said, I believe that I am justified in choosing to focus on panpsychism throughout the rest of this thesis.

In this chapter, I have argued that a commitment to the claim that consciousness cannot be reduced to the purely physical leads us to the acceptance of panpsychism, the theory that consciousness is a fundamental and ubiquitous feature of the universe. I have first explained panpsychism and the main motivation for taking it seriously, namely the Russellian argument that there must be something to matter beyond its spatiotemporal structure that physics reveals to us. Then, some general and common objections to panpsychism have been addressed by spelling out the theory in more detail. Finally, I have argued that radical emergence is impossible and that, if we are anti-physicalists, we must understand consciousness as a weakly emergent feature, meaning that it is already present at the smallest, fundamental level of reality. Thus, we have good reasons to accept panpsychism as a serious theory of consciousness.

3. The Combination Problem and Priority Cosmopsychism

The combination problem, commonly seen as the most pressing difficulty for panpsychism, is the question of how smaller subjects come together to form larger subjects. After explaining the specific form of the problem that will be discussed, I will focus on priority cosmopsychism as discussed by Yujin Nagasawa and Khai Wager, a recent and promising new way of addressing the combination problem. Next, after addressing some general objections to the theory, it will be argued that priority cosmopsychism should be rejected because of its highly problematic ontological commitments. More specifically, I will argue that their ontology cannot accommodate empirical processes relevant to the emergence of minds. Because of this, I will argue that combinatory panpsychism is preferable because it subscribes to a less problematic ontology. Consequently, the way of avoiding the combination problem that was proposed by priority cosmopsychists will turn out to be unsatisfying in virtue of their ontology. The conclusion will be that combinatory panpsychism, while plagued by the combination problem, is still the better route to take.

3.1 Panpsychism and the Combination Problem

The combination problem for panpsychism is the question of how small subjects can come together and bring about more complex, larger subjects. There are other forms of the problem, but I will exclusively focus on this form, as it is plausibly the most vicious version of the worry (Goff forthcoming (a): 136). The combination problem was first raised by William James, who argued that experiences of 'feelings' cannot combine into new, further experiences: "Where the elemental units are supposed to be feelings, the case is in no wise altered. Take a hundred of them, shuffle them and pack them as close together as you can (whatever that may mean); still each remains the same feeling it always was, shut in its own skin, windowless, ignorant of what the other feelings are and mean. There would be a hundred-and-first feeling there, if, when a group or series of such feelings were set up, a consciousness *belonging to the group as such* should emerge. And this 101st feeling would be a totally new fact; the 100 original feelings might, by a curious physical law, be a signal for its *creation*, when they came together; but they would have no substantial identity with it, nor it with them, and one could never deduce the one from the others, or (in any intelligible sense) say that they *evolved* it. [...] The private minds do not agglomerate into a higher compound mind." (James 1890/1950: 160)

This formulation of the combination problem is the *subject-summing problem* (Goff forthcoming (a): 136). Stated differently, the subject-summing problem is the question of how microsubjects, the microphysical subjects of experience, combine to yield macrosubjects, the macroscopic subjects of experience such as humans (Chalmers forthcoming). There is a strong intuition that it is hard to make sense of how this summing happens (Goff forthcoming (a): 136). In principle, for any given group of subjects and any further subject, it is possible for the former to exist without the latter, which means that no group of microsubjects necessitates the existence of a further macrosubject (Chalmers forthcoming). In the case of physical combinations, it is easy to see how smaller objects can combine to create more complex objects, such as in the process of arranging planks of wood to create the physical structure that we call 'ship', but this does not translate to panpsychism. This is a serious problem that needs to be addressed. A combinatory form of

panpsychism, committed to a bottom-up ontology where simplicity combines to yield complexity, needs to make sense of how larger subjects can come about from combinations of smaller subjects.

3.2 Priority Cosmopsychism and the Derivation Problem

There are several ways a panpsychist can try to address the combination problem. Some of them are combinatory forms of panpsychism, meaning that they posit that small subjects combine to form larger subjects, while others are non-combinatorial and are thus not committed to the summation of subjects (Chalmers forthcoming). These will not be addressed in this paper, as I wish to focus exclusively on priority cosmopsychism, a new and promising possible solution that tries to avoid the combination problem at the ontological level. One version of this theory was introduced by Nagasawa and Wager (forthcoming) and relies on the priority monist ontology proposed by Jonathan Schaffer (2010). Priority monism is the position that exactly one *basic* object exists – the cosmos. This does not mean that the cosmos is the only object that exists – that would be *existence* monism. This also does not mean that the cosmos is the total sum of its smallest parts, existing only in virtue of them – that would be *pluralism* or *smallism*. This does mean that the cosmos is a maximal actual concrete object, the ontologically *basic* object prior to all others, meaning that all other objects are its parts or ontologically posterior derivate entities of the cosmos (Schaffer 2010: 33, 66–7). In other words, Schaffer's ontology is a top-down ontology where the whole has priority over its parts (2010: 55). An example Schaffer uses to illustrate this is a circle, whose semicircles are taken to be arbitrary partitions, posterior to the circle as such (2010: 31, 47). A further example he uses is an organism, where the whole organism is taken to be prior to its organs (2010: 47).

Priority cosmopsychism is parallel in structure to priority monism, one exception being that it is concerned with phenomenal properties (Nagasawa and Wager forthcoming: 4), while priority monism is solely concerned with actual concrete objects (Schaffer 2010: 33). Priority cosmopsychists see the cosmic consciousness as the properly basic form of consciousness, understanding all particular conscious subjects as abstractions or derivations of the cosmic consciousness (Nagasawa and Wager forthcoming: 5). Again, this does not mean that only one consciousness exists, which would be existence cosmopsychism, but that the cosmic consciousness is ontologically prior to any particular consciousness (Nagasawa and Wager forthcoming: 5). Proponents of the theory take the whole cosmos as the fundamental level of reality that instantiates phenomenal properties. Smallism, the position that the smallest level of reality is fundamental, is thus denied. Moreover, whether or not such a smallest level exists does not impinge on the main claim of priority cosmopsychism. Since the whole cosmos is seen as fundamental, the priority cosmopsychist is not committed to any claim about whether a smallest level of reality exists. This translates to phenomenal properties as well, given that the cosmic consciousness is regarded as prior to smaller forms of consciousness, so whether or not there is a smallest level of phenomenal properties is irrelevant to the main claims of the theory (Nagasawa and Wager 2010: 8).

Since priority cosmopsychists deny that larger subjects emerge from *combinations* of smaller subjects, they effectively avoid the subject-summing version of the combination problem. Instead, smaller subjects are abstractions or derivations of the cosmic consciousness. Based on this, Nagasawa and Wager explain that priority cosmopsychism replaces the combination problem with the less pressing derivation problem: "How could medium-size individual consciousnesses be derived from the cosmic consciousness?" (forthcoming: 9). They argue that,

if we can show that the consciousness of ordinary individuals can be segmented, then the same can be concluded about the cosmic consciousness (Nagasawa and Wager forthcoming: 10). Since it is conceivable that we can indeed divide the consciousness of ordinary individuals, such as dividing a token visual experience into segments based on colour or regions in space, then we could argue that medium-sized individual consciousnesses are derived from the cosmic consciousness in a similar manner.

Nagasawa and Wager offer further possible answers to the derivation problem by adapting the solutions Schaffer offered to the problem for priority monism. The first way to answer is by arguing that the cosmic consciousness is heterogeneous because it instantiates the property of being heterogeneous, meaning that it instantiates heterogeneity because it follows "a particular path through phenomenal configuration space" (Nagasawa and Wager forthcoming: 11). This is parallel to Schaffer's argument that the physical cosmos, as an integrated whole, is heterogeneous because it instantiates the determinable property of being heterogeneous by tracing a particular curve through physical configuration space (2010: 60). The second way to answer is to say that the cosmic consciousness is heterogeneous because of its regionalised properties, where the "cosmic consciousness is a monadic property which bears a relation of redness to one region and blueness to another region" (Nagasawa and Wager forthcoming: 11). Hence, the cosmic consciousness could bear various relations to regionalised phenomenal properties. This is parallel to Schaffer's proposal that the physical cosmos can be heterogeneous, though monadic, because of regionalised properties, meaning that it bears relations of one kind (e.g. redness) to one region and relations of another kind (e.g. blueness) to another region (2010: 60). The final option is the claim that the cosmic consciousness is heterogeneous because of the regionalised instantiation of phenomenal properties, instantiating one property at one location and another property at another location (Nagasawa and Wager forthcoming: 11), which is parallel to Schaffer's claim that the cosmos might be heterogeneous by instantiating one property here and another property there (Schaffer 2010: 60).

The proposed solutions to the derivation problem that Nagasawa and Wager adapt from Schaffer are deeply confusing. In the case of Schaffer's priority monism, it is easier to make sense of regionalised properties, regionalised 'here-and-there' instantiation and of the claim that the cosmos instantiates the property of being heterogeneous in *physical* configuration space. It is not clear, however, how this translates to the *phenomenal* space postulated by Nagasawa and Wager, or what that space even is. Granted, the authors specifically state that they are proposing a blueprint for priority cosmopsychism and not the full theory, so it is understandable that there are points in the paper that need clarification. Still, until a more detailed account of the proposed solutions is offered, I will refrain from addressing them further.

3.3 Objections to Priority Cosmopsychism

Instead, more general objections to priority cosmopsychism will now be addressed. One such objection comes from E. J. Lowe's critique of Schaffer. One of the main points of Lowe's critique is simply that there is no such object as the cosmos. Schaffer understands the cosmos as consisting of "planets, pebbles, particles, and other proper parts" (Schaffer 2010: 33). What Lowe criticises is the claim that the cosmos is a unified proper (concrete) object. Lowe has no problem with accepting that there might be a 'cosmos' that exists as a *mereological sum* of a bunch of objects, but this is not what Schaffer is arguing for and cannot play the role of the properly basic ontologically prior object (Lowe 2012: 93). The reason why it cannot play that role is because such a mereological sum is *not prior* to its parts, but only exists because its parts exist (Lowe 2012: 93). We can apply the same logic to priority cosmopsychism: if there is no cosmos in

Schaffer's sense, then nothing instantiates the cosmic consciousness posited by the priority cosmopsychist.

An answer that might be offered here is that the cosmos is a unified object rather than a mereological sum precisely *because* of the cosmic consciousness. If we have independent reasons to accept the existence of a cosmic consciousness, then this gives us justification for why we should understand the physical cosmos as a properly basic object as well. These independent reasons for postulating a cosmic consciousness must indeed be strong and convincing. But are they? I think not. It seems that the motivation priority cosmopsychists have for postulating such a consciousness is the need to avoid the combination problem. That does not strike me as a reason strong enough to postulate a radically different 'top-down' ontology, one which is not widely accepted. Moreover, as I will argue, this top-down ontology is deeply problematic, so one might be better of remaining a panpsychist and tackling the combination problem from a bottom-up perspective.

One of the reasons for why I think the top-down ontology of priority cosmopsychism is problematic is the observation that there are processes within the universe that we empirically take to be of a bottom-up nature. The most important example, as it directly pertains to the emergence of minds, is the theory of evolution. During a long span of time, simple life forms change and become more complex. Eventually, once a certain threshold has been passed, these life forms begin to be aware of their environment and, after some more time, of themselves. I think it is safe to say that evolution is a bottom-up process. The question now is this: How does a priority cosmopsychist account for the evolutionary bottom-up process that lead to the existence of conscious beings, given their commitment to the claim that particular subjects are derivations or abstractions of the cosmic consciousness? Alternatively, we could pose the same question for the growth of an organism from conception to adulthood. At some point, the organism will reach the stage of growth when environmental awareness and, eventually, self-awareness appear. There are two ways a priority cosmopsychist could answer this question: a) they could say that bottomup processes can exist even within a top-down ontology, or b) they could claim that the emergence of consciousness via derivation or abstraction is not connected or parallel to physical processes such as evolution and growth. Both answers are problematic.

Regarding answer (a), both the priority monist and the priority cosmopsychist could grant that there are bottom-up processes within a top-down ontology. For instance, particular concrete objects that are derivations or abstractions of the cosmos as a whole could 'loop' back up⁵ and combine to form more complex objects. A priority cosmopsychist might claim that phenomenal properties can 'loop' back up in a similar manner. This would avoid the problem I posed earlier regarding evolution and growth, as it is now easy to make sense of how bottom-up processes can happen within a top-down ontology. However, this poses special problems for the priority cosmopsychist. If small subjects are derivations of the cosmic consciousness that can 'loop' back up and combine to form larger subjects, *the combination problem is not avoided*. Since one of the selling points of priority cosmopsychism is that it replaces the combination problem with the supposedly easier derivation problem, this is not a move that they should take if they want to retain the appeal of their theory.

A bottom-up or combinatory panpsychist, on the other hand, has no problems with the theory of evolution or with growth. If simple objects combine into more complex objects, then simple subjects can combine in a parallel manner to yield complex subjects. Moreover, rather than being a problem for combinatory panpsychism, evolution is fully compatible with the

⁵ Jonathan Schaffer has said in discussion that he believes that this is the correct way to interpret physics.

theory. As Thomas Nagel (2012) argued, panpsychism, or something like it, could explain how the process of evolution results in minds that can rationally understand and explain the world. If consciousness is present at the smallest level of reality, then it is present throughout physical processes such as evolution or growth. Thus, for combinatory panpsychism, such bottom-up processes pose no problem whatsoever.

Regarding answer (b), the priority cosmopsychist can claim that consciousness does not come about from physical processes. These processes should be seen as *diachronic*, meaning that they happen *over* time, while the emergence of consciousness should be understood as *synchronic*, meaning that it happens at a *specific point* in time⁶. Thus the priority cosmopsychist can avoid the worry by arguing that consciousness comes about at an instant in time, while physical processes happen over time. Again, it is fair to restate that Nagasawa and Wager offer a blueprint for priority cosmopsychism rather than a full account. A more complex version of priority cosmopsychism would presumably have ways of dealing with this disconnection between bottom-up physical processes and the top-down derivation of consciousness.

That being said, I simply fail to see how such an account could result in a convincing theory. If the problem of the top-down emergence of subjects who have evolved over time is solved by arguing that consciousness is not tied to physical processes, this can at best be classified as counterintuitive and at worst as going directly against empirical observation. There are good reasons to think of evolution and growth as bottom-up processes and that the evolution of minds is parallel in structure to the gradually increasing complexity of physical systems. A theory that quite openly goes against this parallelism has a large bullet to bite. Hence, priority cosmopsychism has to accept that both top-down and bottom-up processes exist within their

⁶ I wish to thank Philip Goff for pointing out the diachronic/synchronic argument to me and to Yujin Nagasawa for presenting the argument that consciousness might not arise from physical processes in personal correspondence.

ontology, but deny that these processes are related when it comes to consciousness⁷. Nagasawa and Wager admit that this is a strong objection to priority cosmopsychism, but remain optimistic about the prospects of solving it⁸. A combinatory panpsychist is committed only to the claim that bottom-up processes exist and has thus no problem with evolution, growth and the parallelism mentioned earlier. It seems that the combination problem is less of a problem than its derivation counterpart simply because it is not tied to the deeply problematic top-down ontology posited by priority cosmopsychists.

I conclude that bottom-up combinatory panpsychism is the preferable option to priority cosmopsychism. The priority cosmopsychist is committed to a highly problematic ontology that cannot properly account for empirical bottom-up processes relevant to the emergence of consciousness. Their ontology results in either the claim that bottom-up processes exist within a top-down ontology and account for the emergence of larger subjects from smaller subjects, thus failing to avoid the combination problem, or in the claim that consciousness is not a result of physical processes, which commits the priority cosmopsychist to a strange and counterintuitive picture of the world. Thus, while the derivation problem might initially seem to be easier than the combination problem, I have argued that it is not, precisely because of the priority cosmopsychist's ontological commitments.

⁷ It is not clear in what way consciousness is physical in priority cosmopsychism if both the claims that it cannot be reduced to physical facts *and* that there is no parallelism between it and physical processes are true. I wonder if this implies a form of dualism and, consequently, if it inherits all of the problems of dualism.

⁸ I would like to thank Yujin Nagasawa and Khai Wager for discussing this issue with me.

4. Phenomenal Bonding, Information and Panpsychism

In previous chapters, I have argued that phenomenal consciousness cannot be reduced to the purely physical in the strict sense of the term, as well as that it is a fundamental and ubiquitous feature of reality. The result of these considerations is panpsychism. I have then introduced the most pressing objection raised to panpsychism - the combination problem. Priority cosmopsychism was assessed as a possible solution to this problem, but it turned out to be highly problematic. In this chapter, I will argue that a form of information-based panpsychism might better explain how small subjects combine to yield bigger subjects. First, I will introduce the phenomenal bonding solution to the combination problem offered by Philip Goff, who argued that there must be a relation between subjects that enables them to combine into further, more complex subjects. Second, I will argue that information could be used to explain how phenomenal bonding takes place. To do that, I will focus on David Chalmers' dual-aspect view of syntactic information and William Seager's account based on semantic information. Third, I will argue that, if we accept panpsychism, a synthesis of these two views might serve to explain both physical processes and phenomenal bonding. Finally, I will address the objection, raised by Derk Pereboom, that microphenomenal properties cannot ground microphysical properties. I will do so by drawing an analogy to John Heil's ontology of properties, and argue that consciousness does not ground microphysical properties, but that it simply is powerful. The upshot, I believe, is an elegant and explanatorily powerful form of panpsychism, which I call informational panpsychism.

4.1 The Phenomenal Bonding Solution

The phenomenal bonding solution to the combination problem was introduced by Philip Goff (forthcoming (b)). He argues that the argument against the summing of subjects, as discussed earlier, does not exclude that a group of subjects can be in a certain *state of affairs* that can account for the existence of a further subject (Goff forthcoming (b): 8–9). In other words, it is possible that subjects can be *related* in a specific way that necessitates the existence of a further and distinct subject of experience (Goff forthcoming (b): 9). While the no summing of subjects argument rules out that the mere existence of subjects and their intrinsic conscious nature necessitates a further subject with its own distinct phenomenal character (Goff forthcoming (b): 9–10), it does not exclude the following:

"<[S]ubject of experience S1 with phenomenal character x bears relationship R to subject of experience S2 with phenomenal character y> which necessitates <subject of experience S3 exists with phenomenal character z>." (Goff forthcoming (b): 9)

This is phenomenal bonding, a hypothetical relation of how subjects can come together to produce further subjects. As Goff argues, there is no reason to construe panpsychism as a theory that is committed to the claim that the *mere existence* of conscious subjects necessitates a further subject (forthcoming (b): 10). The examples he uses to demonstrate this are organisms and car engines, objects that exist not because of the mere existence of their parts, but because of their parts being related in an appropriate way (Goff forthcoming (b): 10).

There are three reasons why we should take phenomenal bonding seriously. The first one is to defend and make sense of panpsychism, a theoretically attractive and elegant theory that helps explain our complex epistemic situation regarding consciousness (Goff forthcoming (b): 11). The second reason goes beyond our desire to defend panpsychism. As Goff argues, we have reasons to think that there are relations we do not have transparent conceptions of, such as the spatiotemporal relations between objects for which physics only provides mathematical conceptions (forthcoming (b): 12). Microphysics postulates fundamental properties such as mass and charge, but it also postulates fundamental relations, such as spatiotemporal relations. However, as is the case with mass and charge, physics only provides us with the mathematical structure of fundamental relations, but not with their categorical, intrinsic nature (Chalmers forthcoming). As Chalmers argues, it is possible that a phenomenal relation could serve as a categorical basis underlying spatiotemporal relations (forthcoming). The third and final reason is that, analogous to how physics postulates fundamental relations between objects, it is plausible for a panpsychist to assume that there is such a relation that can explain the phenomenal bonding of subjects. Thus, in agreement with Chalmers who sees phenomenal bonding as one of the more promising approaches to the combination problem (forthcoming), I think that we have good reasons to take Goff's proposal seriously.

That being said, we have no transparent conception, one where the nature of the phenomenon is fully understood, regarding the phenomenal bonding relation. If we had such a conception, the combination problem would never have been raised (Goff forthcoming (b): 9). In other words, we would understand precisely how smaller subjects combine into larger subjects, so the solution to the problem would be obvious (Goff forthcoming (b): 9). But we do not have such a conception, so embracing Goff's solution results in a form of mysterianism⁹ about how phenomenal bonding might occur (Goff forthcoming (b): 10). The relations that we can observe empirically are not conceived of as phenomenal bonding relations just as the physical properties

 $^{^{9}}$ In this context, 'mysterianism' roughly means that the concept in question is beyond the grasp of our understanding.

of physics or neuroscience are not conceived of as phenomenal qualities in those sciences (Goff forthcoming (b): 10). However, considering the aforementioned reasons for taking this solution seriously, I think that we should be optimistic about the prospects of explaining it. Perhaps, as Goff argues, we could "identify [phenomenal bonding] with some relation we can observe in the world, or some relation that features in physics" (forthcoming (b): 10).

In the rest of this chapter, I will argue that information could serve as a potential candidate, one that plausibly features in physics, for explaining how phenomenal bonding between subjects might take place. Before doing that, it is important to explain what information is. Colloquially, the term 'information' is used "as an abstract mass-noun to denote any amount of data, code or text that is stored, sent, received or manipulated in any medium" (Adriaans 2012). Explaining information more formally is a difficult task considering the extensive debate centred on its definition. As a result, there is a wide variety of definitions of information. For the purposes of this thesis, I think it is enough to briefly define now and in more detail later the notions of information I will exclusively focus on. The first one is the *formal* or *syntactic* notion of information discussed by Claude Shannon (1948) and used by David Chalmers, who recognised its potential significance for a theory of consciousness (1995; 1996). The syntactic theory of information is concerned with bits or information states that occupy information space. The difference between information states completely accounts for what information is under this theory, rather than the meaning or interpretation of those states. The second one is the *semantic* notion of information, where information is understood as *meaningful data* and not just as a difference between certain information states (Floridi 2015). William Seager (1995) relied on this understanding of information as a possible way of addressing the combination problem.

4.2 Chalmers' Dual-Aspect Theory of Information

I will start by explaining Chalmers' dual-aspect view of information in more detail. He relies on the syntactic or formal notion of information explicated by Shannon (1948). On this understanding, where there is information, there are *information states* that occupy an abstract information space that constitutes the basic structure for a difference relation between those states (1996: 277–8; 2006b). For an example of an information space, take the binary states of 0 and 1. There is a primitive difference between them and the fact that these two states are different completely exhausts their nature (Chalmers 1996: 278). This means that this information space "is fully characterized by its difference structure" (Chalmers 1996: 278). Information spaces are abstract objects, but "following Shannon we can see information as physically embodied when there is a space of distinct physical states, the differences between which can be transmitted down some causal pathway" (Chalmers 2006b: 245). The example Chalmers uses to illustrate how information states are physically embodied in information space is a light-switch, where one state is 'up' and the other is 'down' (1996: 281). A light-switch can take on a potentially infinite number of positions, but the only relevant difference is the one between 'up' and 'down', as only these states make a difference to the light (Chalmers 1996: 281). The switch can be seen as realizing a two-state information space, where one of its physical states corresponds to one information state and another physical state to a different information state (Chalmers (1996: 281).

Based on the observation that there is an isomorphism between physically embodied information spaces and phenomenal information spaces, Chalmers goes on to argue that information space can be realised both physically and phenomenally (2006b: 245). He takes this double-aspect view of information as a possible basic principle, based on which experience arises

as one aspect of information, while the other aspect is embodied in physical processing (Chalmers 2006b: 245). Similar to how information states occupy information space, he argues that states of experience occupy phenomenal information space (Chalmers 1996: 284). For instance, a specific experience of the colour red is one phenomenally realised information state, while the experience of the colour green is another (Chalmers 1996: 284). In the case of phenomenal information space, the difference between phenomenal information states is accounted for by their intrinsic quality and the structure these states bear to each other (Chalmers 1996: 284). Every token experience bears certain similarity and difference relations to other experiences, so that we can always find information spaces into which experiences fall (Chalmers 1996: 284).

However, Chalmers admits that this view is highly speculative. He also states that phenomenal properties seem to have an intrinsic nature over and above their location in information space (1996: 303–4). To make sense of this, we could argue that there is an epistemic gap between the complete index of knowledge about information states occupying information space and phenomenal facts. If that is the case, then it is not clear how a syntactic view of information can accommodate phenomenal consciousness and, consequently, be used to explain phenomenal bonding. It is important to note that Chalmers does not intend to explain consciousness in terms of syntactic information or to propose it as an explanation of phenomenal bonding. I am merely evaluating his view in terms of how it could serve the aim of this chapter: to provide a possible account of phenomenal bonding and thus answer the combination problem. What Chalmers does say which is relevant for the purposes of this chapter is the broader claim that information has both a physical and phenomenal aspect, and that these aspects both pertain to the fundamental level of reality.

4.3 Seager on Semantic Information and Consciousness

I will now explain William Seager's view of information in more detail. Seager (1995) argues for the *semantic* notion of information as a possible solution to the combination problem. Information is semantic when it is not only cast in terms of bit capacity or how much information states or bits an information space can carry, as in Chalmers' view, but when the bits or information states are taken to carry *semantically significant content* (Seager 1995: 277). Put simply, semantic information is *meaningful data* or data *about* something (Floridi 2015). To explain his view, Seager relies on the two-slit experiment where a beam of particles (e.g. photons) is directed towards "an appropriately separated pair of slits in an otherwise opaque surface" (Seager 1995: 277). When the particles hit the detector screen behind the slits, an interference pattern will be formed that results from the interaction of the two paths the particles can take to the screen (Seager 1995: 277). However, the interference pattern disappears if we have a way of determining which path the particles are taking.

To address this, Seager prompts us to consider the case where a quantum eraser is present, a device that erases the information within the detector in order to retrieve the original interference pattern of the particles. As he puts it, the particles now "have no idea whether their paths have been registered or not" (Seager 1995: 278). This strongly suggests that the particles are responsive to the informational state of the quantum eraser (Seager 1995: 278). Furthermore, the experiment would be constructed in a way to prevent the establishment of any sort of causal process between the particles and the quantum eraser. The conclusion Seager draws here is that there is a "noncausal, but *information laden* connection amongst the elements of a quantum system" (Seager 1995: 278). The reason why he favours semantic over syntactic information is because the quantum eraser scenario demonstrates that a syntactic notion, if it is taken to be about

causal processes, is out of the question as there are no causal chains between the eraser and the particles. In other words, the information in question must be semantic because the experiment excludes the possibility of a bit channel or any other form of causal chain exerting influence on the particles passing through the slits (Seager 1995: 278–9).

Based on this, Seager postulates that semantic information might serve to answer the combination problem. We can say that the particles are *informed about* the results of the perfect detector in the two-slit experiment, while they are *informed whether* information has been erased or not in the case of the quantum eraser (Seager 1995: 283). Later on, he argues that this semantic sense of information is a "hidden variable" that exists in addition to the standard physical worldview (Seager 1995: 284). Seager speculates that there is a connection between summation principles, so that when a superposition of states of physical elements happens, there is also mental combination (Seager 1995: 284-5). More specifically, he postulates that quantum coherence might underlie the emergence of complex states of consciousness, as only "coherent multiparticle systems will preserve the peculiar quantum mechanical properties that underlie the appropriate 'summation rules'" (Seager 1995: 285). In the case of quantum coherence, particles are not mere mereological aggregates or mixtures, but combine into "a genuinely new state with properties observably different from the properties of the mixture" (Seager 1995: 284). In other words, quantum wholes are not just the sums of their parts and, as Seager speculates, this might be how subjects combine into distinctive mental wholes (1995: 284). Thus, there is a specific sense in which the semantic notion of information applies to phenomenal bonding and, consequently, to the combination problem.

4.4 A New Dual-Aspect Theory of Information

I will now argue that Seager's semantic information account is potentially a good candidate for Goff's concept of phenomenal bonding, the hypothetical relation of how smaller subjects come together to form bigger subjects (Goff forthcoming (b): 9). Based on the two-slit and quantum eraser experiments, we have good reasons to think that there is a semantically significant information-laden noncausal connection among the elements of a quantum system. The system in question seems to be *informed* by the state of the detector or the eraser. This relation of informing might explain the summing of subjects, perhaps through quantum coherence and superposition, as mentioned earlier (Seager 1995: 285). Even though the details might be unclear, I see semantic information as a promising approach to how phenomenal bonding might take place. In agreement with Seager, who argued that the notion of a system *being informed* could be seen as a part of the realm of mentality (1995: 283), I believe that semantic information is a natural correlate to consciousness.

As Wheeler (1990) and Fredkin (1990) have argued, physical relations might ultimately be explained in terms of information. Referenced by Chalmers as well (1996: 302), this idea is based on the syntactic or formal notion of information. However, as I have argued, consciousness cannot be reduced to the purely physical even if the physical is cast in terms of syntactic information, as there is an epistemic gap between the complete index of knowledge about information states occupying information space and phenomenal facts. So if consciousness cannot be reduced in this way, a panpsychist needs a way of making sense of how subjects combine which is not cast in terms of syntactic information. Therefore, I argue that we might be able to explain how phenomenal bonding happens with reference to semantic information. The semantic notion is still a notion of *information*, so it is plausibly more parsimonious to postulate it instead of a significantly different kind of notion. If we end up with a picture of the world where physical relations and the phenomenal bonding relation can both be explained in virtue of information, this picture will ultimately be more elegant than those that postulate differing notions.

Information on the view I am proposing has two aspects, one syntactic, which is related to physical relations, while the other is semantic and related to the phenomenal bonding relation. Similar to how Chalmers argues for a double-aspect view of *syntactic* information, where one aspect is physical, the other phenomenal, I instead propose a double-aspect view of *information*, where one aspect is syntactic and the other semantic. If consciousness cannot be reduced to syntactic information, the introduction of semantic information to serve as the phenomenal aspect of a double-aspect view, as well as to explain how phenomenal bonding might take place, seems to be an attractive proposal. To recap, this is my complete view:

1) There is a fundamental level of reality consisting of the smallest physical ultimates,

2) The intrinsic nature of these physical ultimates is consciousness,

3) The physical relations between these ultimates can be explained in terms of syntactic information,

4) The phenomenal bonding of simpler subjects into more complex subjects can be explained in terms of semantic information.

Some systems will exhibit a structure of the smallest physical ultimates where phenomenal bonding occurs between them, meaning that this system will result in a larger subject. The human

brain could be taken as an example of such a system. Other systems will not result in phenomenal bonding, meaning that the structure of physical ultimates will be a mere heap or aggregate of the smallest subjects, with no resulting larger subject. A rock or a table could, presumably, be taken as examples of such a structure. Goff (forthcoming (b)) claims that we do not have a transparent conception of phenomenal bonding. I have made a proposal for what phenomenal bonding might be, so I believe that this is an improvement or at least a good first step towards a more complex account.

4.5 Consciousness and Grounding

I believe that (2) requires further explanation. It is important to clarify what I mean when I say that consciousness is the intrinsic nature of physical ultimates, so as to avoid an important problem posed by Derk Pereboom (2015). He argues that: "the envisioned phenomenal micropsychism¹⁰ would need to posit brute laws linking microphenomenal absolutely intrinsic properties with the conventional microphysical properties that they underlie, without which the truths about the microphysical properties would not be derivable from the micropsychist truths" (Pereboom 2015: 317). Put simply, if consciousness is the intrinsic nature of matter that grounds (micro)physical properties, there need to exist metaphysically brute laws, meaning that they are primitive and that they cannot be explained any further, to link consciousness (the microphenomenal) with the (micro)physical properties it grounds.

However, on my view, consciousness does not *ground* physical properties. Consciousness simply *is* causally efficacious. In other words, consciousness should be understood as a powerful quality, something which is at the same time dispositional and categorical. This prompts

¹⁰ It is important to note that Pereboom uses the term 'micropsychism' here to denote the theory that phenomenal properties are present at the smallest level of reality. This is different from Strawson's notion of micropsychism used earlier in this paper, which denotes the view that only *some* physical ultimates instantiate phenomenal properties, but not all.

clarification. Dispositional properties are the powers or dispositions of an object for particular kinds of manifestation (Heil 2003: 11). It is important to note that the microphysical properties discussed by Pereboom simply *are* dispositional properties, that is, they are characterised by physics in terms of their dispositions. Categorical properties are the properties of objects that serve to ground dispositional properties (Heil 2003: 29). The way I understand consciousness as a powerful quality should be understood in analogy to John Heil's ontology of properties. Traditionally, the discourse around properties revolves around two possible options: 1) all properties are dispositional properties, or 2) dispositional properties are grounded in categorical properties (Heil 2003: 29). Rather than arguing for one of those views, Heil argues for an identity between dispositional and categorical properties:

"If *P* is an intrinsic property of a concrete object, *P* is simultaneously dispositional and qualitative; *P*'s dispositionality and qualitativity are not aspects or properties of *P*; *P*'s dispositionality, P_d , is *P*'s qualitativity, P_q , and each of these is $P:P_d = P_q = P$." (2003: 111)

This means that every property of an object is at the same time dispositional and categorical. For example, the sphericity of a ball is a quality possessed by the ball and is at the same time its power (e.g. to roll) (Heil 2003: 112). In a similar manner, I argue that consciousness is categorical and at the same time powerful. It is categorical as the intrinsic nature of physical ultimates and powerful via the syntactic and semantic aspects of information, which explain how physical and phenomenal relations between these ultimates might take place. If we can identify the conscious categorical properties with the dispositional properties described by physics, then we have accounted for the microphysical properties discussed by Pereboom. This is a strange suggestion, but I think that it does offer a way of avoiding Pereboom's objection: consciousness

does not *ground* microphysical properties, consciousness simply *is* powerful, that is – a powerful quality.

Admittedly, what I have argued for is highly speculative. It is unclear how to define information, and even less so when we are dealing with more specific accounts of information. However, if physical relations can potentially be cast syntactically, in terms of primitive differences between information states occupying information space, I believe that postulating semantic information, as a way of explaining how the phenomenal bonding relation might take place, is at least parsimonious. The result is a dual-aspect view of information, where information is taken to have a syntactic and semantic aspect. Thus, I tried to explicate both physical relations and the phenomenal bonding relation through a singular notion: information. On this informational form of panpsychism, physical ultimates are taken to instantiate phenomenal properties which are at the same time categorical and dispositional. Small subjects phenomenally bond together to form further subjects when they stand in an appropriate semantic-informational relation. Because of the elegance of this theory, I think it should be taken seriously as a blueprint for a more detailed form of informational panpsychism. Moreover, since physical relations are understood in terms of information as well, informational panpsychism might serve not only to explain the nature of consciousness and of phenomenal bonding, but as a complete theory of the universe.

Concluding Remarks

In this thesis, I have first argued that consciousness cannot be explained in purely physical terms. I have done so by arguing that there is an extreme epistemic gap between our knowledge of physical and phenomenal facts. Because of this, we have good reasons to reject physicalism. If this conclusion is accepted, it leads us to the acceptance of panpsychism, the theory that consciousness is a fundamental and ubiquitous feature of reality. This was explained through a discussion of emergence. I have argued against strong or radical emergence, showing that it is highly improbable. What follows from this is that we should take consciousness as a weakly emergent feature. If consciousness comes about from combinations of smaller physical ultimates, then those ultimates must already instantiate consciousness. Otherwise, consciousness would be radically emergent, which is impossible. To strengthen the claim of the impossibility of radical emergence, I emphasised the point that no cases of it have ever been observed. If this is taken seriously, panpsychism naturally follows.

I have then focused on the combination problem: how do small subjects or microsubjects combine to yield larger subjects or macrosubjects? After explaining this problem in greater detail, I introduced priority cosmopsychism as a possible way of avoiding it. However, because of the deeply problematic ontological commitments of priority cosmopsychism, I have argued that we should reject it and instead remain panpsychists. Finally, through a discussion of David Chalmers and William Seager, I have suggested that physical relations and Philip Goff's phenomenal bonding relation can be explained in terms of syntactic and semantic information, respectively. The result of these considerations is informational panpsychism, an elegant theory that could, if I may speculate, possibly serve as a new paradigm for how we understand the universe.

References

Adriaans, P. (2012) 'Information', *Stanford Encyclopedia of Philosophy*. <URL= http://plato.stanford.edu/entries/information/>.

Brentano, F. (1973) Psychology from an Empirical Standpoint. London: Routledge.

Chalmers, D. (1995) 'The Puzzle of Conscious Experience', reprinted from *Scientific American* 273: 80–86.

 $<\!\!\text{URL=http://www.rudygarns.com/class/mind/lib/exe/fetch.php/chalmers.puzzle_conscious_experience.pdf\!>.$

Chalmers, D. (1996) *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.

Chalmers, D. (2003) 'Consciousness and its Place in Nature', in S. P. Stich and A. Warfield (eds.) *The Blackwell Guide to the Philosophy of Mind*, 102–142. Oxford: Blackwell Publishing.

Chalmers, D. (2006a) 'Strong and weak emergence', in P. Davies and P. Clayton (eds.) *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion*, 244–254. Oxford: Oxford University Press.

Chalmers, D. (2006b) 'Facing up to the Problem of Consciousness', in Maureen Eckert (ed.) *Theories of Mind: An Introductory Reader*, 224–249. Lanham, Maryland: Rowman & Littlefield Publishers, Inc.

Chalmers, D. (2007) 'Phenomenal concepts and the explanatory gap', in T. Alter and S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, 167–194. Oxford: Oxford University Press.

Chalmers, D. (2015) 'Panpsychism and Panprotopsychism', in T. Alter and Y. Nagasawa (eds.) *Consciousness in the Physical World: Essays on Russellian Monism*, 246–276. Oxford: Oxford University Press.

Chalmers, D. (forthcoming) 'The Combination Problem for Panpsychism', in G. Bruntrup and L. Jaskolla (eds.) *Panpsychism*. Oxford: Oxford University Press.

Díaz-León, E. (2010) 'Can Phenomenal Concepts Explain The Epistemic Gap?', *Mind* 119 (476): 933–51.

Floridi, L. (2015) 'Semantic Conceptions of Information', *Stanford Encyclopedia of Philosophy*. <URL= http://plato.stanford.edu/entries/information-semantic/>.

Fredkin, E. (1990) 'Digital mechanics', Physica D 45: 254-70.

Goff, P. (forthcoming (a)) *Consciousness and Fundamental Reality*. Oxford: Oxford University Press.

Goff, P. (forthcoming (b)) 'The Phenomenal Bonding Solution to the Combination Problem', in G. Bruntrup and L. Jaskolla (eds.) *Panpsychism*. Oxford: Oxford University Press.

Heil, J. (2003) From an Ontological Point of View. Oxford: Clarendon Press.

Jackson, F. (1982) 'Epiphenomenal Qualia', Philosophical Quarterly 32: 127-136.

James, W. (1890/1950) The Principles of Psychology. New York: Dover Publications, Inc.

Lowe, E. J. (2012) 'Against Monism', in Philip Goff (ed.) Spinoza on Monism, 92–112. Basingstoke: Palgrave Macmillan.

Nagasawa, Y., Wager, K. (forthcoming) 'Panpsychism and Priority Cosmopsychism', in G. Bruntrup and L. Jaskolla (eds.) *Panpsychism*. Oxford: Oxford University Press.

Nagel, T. (1974) 'What is it like to be a bat?', Philosophical Review 83: 435-50.

Nagel, T. (1979) Mortal Questions. Cambridge: Cambridge University Press.

Nagel, T. (2012) *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature is Almost Certainly False.* Oxford: Oxford University Press.

Nida-Rümelin, M. (2009) 'Qualia: The Knowledge Argument', *Stanford Encyclopedia of Philosophy*. <URL=http://plato.stanford.edu/entries/qualia-knowledge/>.

Papineau, D. (2007) 'Phenomenal and Perceptual Concepts', in T. Alter and S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, 111–144. Oxford: Oxford University Press.

Pereboom, D. (2015) 'Consciousness, Physicalism, and Absolutely Intrinsic Properties', in T. Alter and S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, 300–323. Oxford: Oxford University Press.

Russell, B. (1927/1992) The Analysis of Matter. London: Routledge.

Schaffer, J. (2010) 'Monism: The Priority of the Whole', *Philosophical Review*, 119 (1): 31–76.

Seager, W. (1995) 'Consciousness, Information and Panpsychism', *Journal of Consciousness Studies*, 2 (3): 272–88.

Seager, W. (2006) 'The 'Intrinsic Nature' Argument for Panpsychism', *Journal of Consciousness Studies*, 13 (10–11): 129–145.

Seager, W. (2010) 'Panpsychism', *Stanford Encyclopedia of Philosophy*. <URL=http://plato.stanford.edu/entries/panpsychism/>.

Shannon, C. E. (1948) 'A mathematical theory of communication', *Bell Systems Technical Journal* 27: 379–423.

Stoljar, D. (2005) 'Physicalism and Phenomenal Concepts', Mind & Language 20 (5): 469-494.

Strawson, G. (2006) 'Realistic Monism, Why Physicalism Entails Panpsychism', *Journal of Consciousness Studies*, 13 (10–11): 3–31.

Van Gulick, R. (2014) 'Consciousness', *Stanford Encyclopedia of Philosophy*. <URL=http://plato.stanford.edu/entries/consciousness/>.

Vintiadis, E. (No date) 'Emergence', *The Internet Encyclopedia of Philosophy*, ISSN 2161-0002. <URL=http://www.iep.utm.edu/emergenc/>.

Wheeler, J.A. (1990) 'Information, physics, quantum: The search for links', in W. H. Zurek (ed.) *Complexity, Entropy, and the Physics of Information*. Redwood City, CA: Addison-Wesley.

Whitehead, A. N. (1967) Adventures of Ideas. New York: The Free Press.