

This work is protected by copyright and other intellectual property rights and duplication or sale of all or part is not permitted, except that material may be duplicated by you for research, private study, criticism/review or educational purposes. Electronic or print copies are for your own personal, non-commercial use and shall not be passed to any other individual. No quotation may be published without proper acknowledgement. For any other use, or to quote extensively from the work, permission must be obtained from the copyright holder/s.

# **From quarks to cognition: in pursuit of a complete micropsychist theory of consciousness**

**Benjamin James Hibbert**

Thesis submitted for the degree of Doctor of Philosophy in Philosophy

June 2022

Keele University

This thesis could not have happened without the support of my loved ones, who know who they are. They are the rocks that have enabled this passage through the river of time, and I remain eternally grateful for their support and love.

*There is nothing that we know more intimately than conscious experience, but there is nothing that is harder to explain'*

(Chalmers 1995, p. 200)

## *Abstract*

Despite having received a recent resurgence of interest, Panpsychism remains a worldview afflicted by various explanatory shortfalls. Chiefly, the literature is yet to provide a ‘proper’ panpsychist theory of consciousness with the potential to explain precisely how we have transitioned from a myriad of rudimentary, entropic instantiations of phenomenality to unified and bounded macro-level conscious subjects replete with an awareness of their ontology. This work recognises this particular shortfall and attempts to remedy it by offering a ‘complete micropsychist theory of consciousness’ that is capable of overcoming the various iterations of the combination problem, and explaining how the sort of minimal phenomenality associated with bottom-level physical entities might be reconciled with the sort of cognitive phenomenology we undergo as higher-level subjects, by advancing a new take on the nature of what a phenomenal property is that might very well lead us to both reanalyse our understanding of the natural world and rethink our own nature.

# Table of Contents

<b>Chapter 1</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>1</b>
<b>1.1 Constitutive Panpsychism and its place within contemporary philosophy of mind.</b> .....	<b>9</b>
<b>1.2 Why constitutive Russellian micropsychism is (potentially) the most explanatorily complete form of panpsychist Russellian monism</b> .....	<b>15</b>
1.2.1 Panprotopsyichism.....	17
1.2.2 Non-constitutive Panpsychism .....	21
1.2.3 Priority Cosmopsychism.....	28
<b>1.3 The fundamental problems facing Constitutive Micropsychism</b> .....	<b>42</b>
1.3.1 The Combination Problem(s): Making sense of subjects constituting other conscious subjects .....	44
1.3.2 The ‘Inverted’ Combination Problem: reconciling Micropsychism with our intuitions .....	56
<b>Chapter 2</b> .....	<b>60</b>
<b>On the nature of phenomenal properties</b> .....	<b>60</b>
<b>2.1 In favour of Powerful Qualities: On the incompleteness of the physical account of causation and the causal relevance of phenomenal properties</b> .....	<b>64</b>
2.1.1 The Self-Stultifying Objection to Epiphenomenalism .....	66
2.1.2 Expanding the concept of causation .....	76
<b>2.2 “Experiences-for”: On subjects of experience and the impossibility of unexperienced instantiations of phenomenality.</b> .....	<b>86</b>
2.2.1 Subjects, Subjects of experience and Subjectivity .....	99
2.2.2 Minimal Subjects of Experience.....	108
<b>2.3. “Experiences-of”: On experiential content and the weak transparency of phenomenal properties.</b> .....	<b>115</b>
2.3.1 A problem for representationalist accounts of phenomenal content .....	117
2.3.2 Three ways <i>not</i> to overcome the problem.....	127
2.3.3 Towards a potential solution.....	134
<b>2.4. Plugging the hole in the physical</b> .....	<b>139</b>
<b>Chapter 3</b> .....	<b>144</b>
<b>A Place for Phenomenal Properties</b> .....	<b>144</b>
<b>3.1 Motivating Liberal Naturalism</b> .....	<b>146</b>
3.1.1 Naturalism and the problem of ontological plurality.....	148
3.1.2 Liberal Naturalism: A solution to the problem.....	152
<b>3.2 Motivating the move to reject bottom-level categorical phenomenal properties</b> .....	<b>159</b>
<b>3.3 Powerful Phenomenal Qualities: Towards a liberally naturalized ontology of phenomenality</b> .....	<b>164</b>
3.3.1 On the nature of powerful qualities .....	165
3.3.2 Russellian Panpsychism and Identity theories of powers.....	167
3.3.3 Towards a bundled theory of fundamental phenomenal powers .....	174
3.3.4 Where does this leave us?.....	179
<b>Chapter 4</b> .....	<b>183</b>
<b>A Role for Phenomenal Properties</b> .....	<b>183</b>
<b>4.1 On phenomenal properties and Individuation</b> .....	<b>185</b>

4.1.1 Grounding an account of first-order phenomenal knowledge .....	188
4.1.2 Phenomenal Continuity and Phenomenal Overlap .....	192
4.1.3 Reframing Phenomenal individuation .....	195
<b>4.2 Towards an account of Phenomenal Maxwellian Demons .....</b>	<b>198</b>
4.2.1 Phenomenal Properties realizing Phenomenal Maxwellian Demons .....	202
4.2.2 On the possibility of P-Demons.....	211
4.2.3 A place for Phenomenal Individuation and P-Demons. ....	214
4.2.4 The Constructal Law ‘from the inside’ .....	224
<b>Chapter 5 .....</b>	<b>232</b>
<b><i>Addressing the ‘first-order’ combination problems .....</i></b>	<b>232</b>
<b>5.1 Towards a solution to the first-order combination problems.....</b>	<b>234</b>
5.1.1 A theory of composition as identity .....	236
<b>5.2 Two potential solutions to the combination problem(s) .....</b>	<b>241</b>
5.2.1 Dominant Fusionism.....	244
5.2.2 Identity Panpsychism and Dominant Phenomenality .....	248
5.2.3 Non-fused Dominant Phenomenal Overlap.....	252
<b>5.3 A consequence of these solutions: A response to the palette problem. ....</b>	<b>261</b>
<b>Chapter 6 .....</b>	<b>265</b>
<b><i>Towards a Complete Micropsychist Theory of Consciousness .....</i></b>	<b>265</b>
<b>6.1 The evolution of P-Demons: A solution to the second-order combination problems .....</b>	<b>271</b>
6.1.1 Informational Entropy and Phenomenal Surprise.....	272
6.1.2 Amplifying Phenomenal Information.....	274
6.1.3 An insight into the Structural Mismatch Problem .....	280
<b>6.2 The Phenomenality of Thought .....</b>	<b>287</b>
6.2.1 A new take on Cognitive Phenomenology .....	288
6.2.2 Addressing the Awareness Problem .....	296
<b>Chapter 7 .....</b>	<b>300</b>
<b>Conclusion.....</b>	<b>300</b>

# Chapter 1

## Introduction

Imagine that every instantiation of the macro-level properties we bear witness to in the universe could be reduced to the micro-level properties contained within a set of ‘infinity dice’ (dice with an infinite number of possible sides that exemplify a potentially infinite number of possible states the universe may occupy)<sup>1</sup>. Now, imagine that the experience you are currently having is *not* illusory, strong emergence<sup>2</sup> *is* impossible, and that the physicalist is well-motivated in her commitment to the contention that such dice must exclusively contain *narrowly* physical properties<sup>3</sup> that are exclusively structural and wholly non-experiential. Given these conditions, one might ask how, upon rolling these purely physical ‘infinity dice’ infinitely, could we ever account for the phenomenal truth that there *is* ‘something it is like’ (Nagel 1974, p. 442) for a conscious subject to bear witness to the blackness of these letters contrasted against the white backdrop, or muse upon such lofty concepts as dice with infinite sides? Rolling this hypothetical infinity dice therefore invites

---

<sup>1</sup> I note that the use of dice is in this case somewhat arbitrary. We might supplant ‘infinity dice’ for any fundamental entity, or bundle of properties, with a potentially infinite quantity of distinct possible states, or ‘resting positions’, that might be brought about by, deduced from, or reduced to these fundamental starting conditions (i.e., in the case of dice, an infinite number of distinct possible states that this entity might bring about or occupy upon being ‘rolled’). The point to be laboured here is that regardless of the entity employed, if this entity, or set of fundamental properties, is exclusively narrowly physical in such a way that does not contain, or does not necessitate the ontological possibility of ever achieving, non-illusory phenomenality, the ‘dice’ will only contain the potential for an infinite number of resting positions of a certain *narrowly physical* kind and thereby there will never be *any* resting position that is reconcilable with a non-illusory account of the phenomenal experiences we currently appear to be having.

<sup>2</sup> See p.11 for an account of what is meant by strong emergence in this context.

<sup>3</sup> In line with Goff (2017) and Chalmers (2016), I define *narrowly* physical properties as the purely spatiotemporally extended, or structural properties, of physical entities, which are wholly non-phenomenal and non-protophenomenal in nature. I note here that physicalism, in its strongest articulation, maintains that all facts are *narrowly* physical facts, or are constituted by said *narrowly* physical facts. I also note that this concept might be contrasted with *broad* physical properties and facts, which are instantiations of physicality that are *not* exhausted by an account of structural descriptions.



the question: from where amidst this conglomeration of microphysical truths could we ever account for experiential truths?

As a result of this ‘hard’ problem, and others of a not dissimilar ilk<sup>4</sup>, Panpsychism has seen a resurgence of interest within contemporary philosophy of mind (with advocates including Strawson 2006, Seager 2006, Nagel 2012, Basile 2010, Goff 2017, Chalmers 1996), and promises to solve such problems by positing phenomenal properties as ontologically fundamental constituents of the universe. The most popular form of panpsychism is a ‘bottom-up’<sup>5</sup>, or ‘Smallest’ Constitutive Panpsychism (or Constitutive Micropsychism), which holds that 1) the universe is fundamentally microphysical, 2) the nature of the microphysical is ubiquitously extrinsically spatial *and* intrinsically experiential, and 3) all macro-level experiential facts we bear witness to are constituted by more fundamental and rudimentary micro-level experiential facts. Advocates maintain that, given the myriad of unsavoury consequences involved in endorsing competing metaphysical commitments, and given the seemingly irreconcilable dichotomy between *narrow* physicalism and our existence as subjects of experience, a constitutive panpsychism of this kind simply ‘must be true’ (Basile 2010, p. 93). However, whilst it is easy to see how this metaphysic may avoid the ‘hard’ problem of explaining how experiential properties arose from wholly non-experiential properties, the explanatory success of constitutive panpsychism in relation to the ‘hard’ problem is hampered by a lack of explanatory success in three key areas. First, the

---

<sup>4</sup> I note here that this particular ‘hard’ problem is subtly distinct from the more widely endorsed formulation advanced in Chalmers (1996). I also note, however, that both Chalmers’ formulation and my own are similar in so far as they highlight the ontological and explanatory gap that occurs upon attempting to reconcile phenomenal properties with narrow physicalism.

<sup>5</sup> This term is borrowed from Goff (2017), who distinguishes between a ‘bottom-up’ constitutive micropsychism and a ‘top-down’ constitutive cosmopsychism. The former posits experiential facts as grounded in more rudimentary experiential facts that exist at a micro-level, such that there is ‘something it is like’ for a quark or an electron. Whereas the latter posits experiential facts as grounded in experiential facts concerning the universe, such that there is ‘something it is like’ for the macro-level cosmos. This distinction shall be expanded upon in section 1.2.3.

constitutive panpsychist faces the issue of establishing precisely how such ubiquitous experiential properties may be reconciled with our intuition that certain physical substances, such as thermostats, rocks, or chairs, are not experientially constituted (avoiding this problem, as I shall articulate in 1.3.2, involves either resigning panpsychism to the counterintuitive claim that rocks have experiences, or advancing a solution to, what I term, the ‘inverted’ combination problem). Further, panpsychism faces the more standard ‘combination problem(s)’ (see James 1890 and Chalmers 2016) of explaining how a variety of distinct experiential properties replete with individually rudimentary phenomenal experiences simultaneously constitute and combine into a singular complexly structured<sup>6</sup> subject of experience with a distinctly unified, experientially varied, subjective quality of ‘what it is like’ (Nagel 1974, p. 442)<sup>7</sup>. Finally, and arguably most importantly, panpsychism is simply a meta-theory of mind that maintains that phenomenal properties are ubiquitous but offers no explanation for precisely what these properties are, precisely what role they play in the natural world, or precisely how these properties produce consciousness as we know it (Skrbina 2007, p. 249). This problem has given rise to panpsychism being discounted as a

---

<sup>6</sup> I make use of the term ‘structure’ throughout this thesis. I note that, in all cases, my use of ‘structure’ is employed to capture the ‘form’ an instantiation of physicality occupies. I take it that my usage is largely uncontroversial and amounts merely to the suggestion that *every material object, concrete entity, or relational bundle of properties, occupies at the very least a physical ‘point’ in space-time and the arrangement of the space that it occupies is its ‘structure’, or perhaps its physical ‘shape’, or ‘form’, which is grounded in the properties it holds* (see Dipert 1997 for a similar account of the term ‘structure’ and see Skow 2007 for a review of the uncontroversial nature of the claim that all physical things are, at the very least, ‘shaped’). On this liberal use of the term, the concept of a structure might be broadly applied to all concrete entities and might be interchangeably used with ‘spatial arrangement’, ‘physical state’, ‘form’, or even ‘shape’, when caveated that all such concepts, in some sense, capture what it is for a concrete instantiation of physicality to occupy a particular physical mode of being in space. I also note that ‘shape’ or ‘structure’ in this context might very well be an intrinsic property (see Lewis 1983), but I do not see any need to explicitly defend this claim here. For my purposes, it suffices to say that every concrete physical entity at least has a spatial structure and for now merely mention that these ‘structures’ might both be used to demarcate between physical objects, and, as shall be revealed in chapter 2, might very well have a special relationship with the content of phenomenal experiences. I also note that, in the later chapters, I employ the use of ‘complex’ or ‘internal’ structure to capture a type of structure that has acquired a substructure, or various parts, and this might be contrasted with other fundamental structures that occupy a physical form which cannot reasonably be described as having parts.

<sup>7</sup> As shall be explicated within section 1.3.1, since James’ (1890) earliest formation of the combination problem for panpsychism, Chalmers (2016) has expanded the issue to one of explaining the combination of *subject, quality and structure*.

‘superficially attractive idea’ that struggles when tasked with performing ‘any sort of explanatory work’ (Humphrey 1992, p. 203), and certain philosophers have highlighted that the only way to avoid this charge of explanatory deficiency is to advance a, as yet unarticulated, ‘proper panpsychist theory of consciousness’ (Chalmers 2016, p. 27) that will explain the nature of micro-phenomenal properties in a manner that also explains how such properties combine into the type of conscious subjects that query their ontology<sup>8</sup>.

These explanatory deficiencies have caused several contemporary philosophers to reject panpsychism outright (See Searle 2004, McGinn 2006, Dennett 1991), and have left many advocates of constitutive panpsychism asking, ‘it must be true – but how can it be?’ (Basile 2010, p. 93). Answering this question shall be the focus of this thesis. I endeavour to offer a complete ‘bottom-up’ constitutive panpsychist theory of consciousness that solves the three aforementioned problems simultaneously, and maintain that *if* some of the most pressing problems associated with panpsychism can be convincingly solved within this model then I will have gone some way to positioning constitutive panpsychism as an explanatorily coherent and attractive worldview that stands as a wholly viable alternative to the physicalism that has dominated much of the contemporary philosophical and scientific zeitgeist. In this respect, this thesis is exclusively concerned with addressing the question: ‘*If* phenomenal consciousness is *not* illusory<sup>9</sup> or strongly emergent, how does it fit into our naturalised worldview in a way that might emancipate constitutive micropsychism from its explanatory shortfalls and explain our existence as macro-level conscious subjects?’. I do not devote this thesis to a detailed examination of the merits of panpsychism over its metaphysical counterparts, nor do I intend to explicitly address any other conceptually or

---

<sup>8</sup> Skrbina (2007) also argues for the need to provide a ‘panpsychist theory of mind’ (p. 249).

<sup>9</sup> I endorse Frankish’s (2016) contention that *if* narrow physicalism is true, then phenomenological consciousness is necessarily illusory.

ontologically motivated solutions to the hard problem. Instead, my focus is fixed with rigidity upon articulating an explanatorily robust panpsychist narrative that simultaneously addresses all iterations of the ‘combination problem’ (see Chalmers 2016 and James 1890) afflicting contemporary ‘bottom-up’ versions of constitutive panpsychism, whilst also explaining how micro-phenomenal properties have transitioned from entropic, rudimentary constituents of the universe to fully-realized subjects of experience replete with a conscious awareness of what appear to be their constituent parts. I maintain that such a narrative would prove to be of significant import to contemporary panpsychism, especially as a number of panpsychists are now calling to reject ‘bottom up’ panpsychist theories in favour of an ostensibly more explanatorily robust ‘top-down’ cosmopsychism (see Goff 2017, Wager & Nagasawa 2016, Shani 2015), and certain once ardent proponents of a ‘bottom-up’ version of constitutive panpsychism have rejected the metaphysic all together on the grounds that the standard combination problem appears insolvable (see Coleman 2009 - 2013). In this sense, this thesis is devoted, in entirety, to repositioning bottom-up constitutive panpsychism as an attractive option within the field of the metaphysics of the mind via articulating a coherently unified constitutive micropsychist theory of consciousness the likes of which has, thus far, been sorely missing from the panpsychist literature.

With this established, I note that the overarching narrative contained herein may be disentangled into six key sub-narratives contained within each of the six primary chapters. The first chapter is concerned with motivating Russellian micropsychism over its (panpsychist) counterparts. The second provides an attempt to earnestly define phenomenal properties from the perspective of Russellian micropsychism. The third and fourth chapters aim to integrate phenomenal experience, so defined, into a naturalised account of reality and find a role for phenomenal properties by accounting for the dispositions that might be imbued

as a result of this naturalised account of phenomenality. The fifth chapter is an attempt to build upon this account to solve (some) formations of the combination problem(s), and the sixth is an attempt to combine all the above to explain how we have transitioned from rudimentarily experiential quarks (or whichever entity turns out to be fundamental) to human, conscious subjects replete with the capacity for cognition<sup>10</sup>. In this sense, this thesis might be described as a narrative of two overarching halves. The former advances reasons for adopting a specific account of bottom-level experiences, whilst the latter builds upon the account of microexperiential properties offered and attempts to construct a model capable of explaining precisely how the qualities and dispositions associated with the experience of ‘what it is like’ to be a quark, or any other fundamental entity, may be naturalized in a manner that makes no observable (or predictive) difference to our scientific worldview but may be exercised so as to provide us with a satisfying resolution to the combination problems and an equally satisfying answer to the fundamental question: *what are we?*

In order to reach this point, I end the first ‘half’ of this overarching narrative by offering reasons to adopt a specific account of phenomenal properties, in which, at the bottom-level, they are identified as powerful qualities that manifest as unitary, rudimentary, physically constrained instantiations of ‘experiences-of-content’ and ‘experiences-for-subjects’ which, when appropriately structured, and standing in the right reciprocal relation to their environment, hold the disposition to realise all of the manifestations documented in physics *and* minimally phenomenally individuate between representational types. With this as grounding, I begin the second ‘half’ of this narrative by attempting to show how a rudimentary theory of the mind predicated upon a form of micropsychism is perhaps the sort

---

<sup>10</sup> I note that such an account might hold the potential to at least offer some insight into the other combination problems that were not addressed in chapter 4 (what I call the second-order combinations problems, see chapters 5/6).

of theory that Maxwell (1871) was attempting to elucidate in order to ground his formation of a micro-level, conscious, negentropic sorting demon<sup>11</sup>, and I argue that upon properly defining microphenomenal properties our best means of conceptualizing a potential physical manifestation of the disposition to phenomenally individuate is to think of this behaviour in terms of naturalised ‘Maxwellian Demons’. From here, I attempt to construct a naturalized model for a ‘Phenomenal Maxwell’s Demon’<sup>12</sup> (a ‘P-Demon’) in the hope of using this to solve some of the recent formulations of the combination problems that so afflict contemporary panpsychism (see Chalmers 2016). In order to achieve this, I endeavour to first move the discussion away from contemporary solutions to the subject-summing problem, which attempt to combine the experiential subject ‘K’ with the experiential subject ‘S’ so as to form a new subject ‘X’, and instead maintain that a solution to this particular problem (and its counterpart – see section 1.3.2) may be found within an account of dominant phenomenal bonding, which adequately resolves this issue but does *not* necessarily require for distinct experiential subjects to be bound together to form a novel subject. From here I reiterate that, once the subject-summing problem has been resolved, the causal role of phenomenal properties is rendered most explicit when thought of in terms of Naturalised Maxwellian Demons, and, emboldened by the work of contemporary physicists who have attempted to move Maxwellian Demons away from their purely theoretical origins (see Busby & Howard 2017 and Bejan 2014), I elucidate a practical model for a ‘Phenomenal Maxwellian Demon’ that would fit neatly into our naturalised scientific worldview in a manner that might solve the hard problem, solve (some of) the combination problems, and also potentially lead the

---

<sup>11</sup> Or is, at least, perhaps not at odds with one specific account of how a substantially less sophisticated Maxwellian Demon might operate in the world.

<sup>12</sup> As shall be articulated in detail in chapter 4, by this I mean a model of a Maxwell’s Demon that does *not* violate Landauer’s (1961) Principle.

way to a true panpsychist theory of consciousness with the potential to explain why it is that ‘we all behave like Maxwell’s Demon’ (Gleick 2011, p. 281).

However, prior to elucidating such a theory, it seems prudent to first address the metaphysical commitments underlying constitutive panpsychism and address some of the contemporary problems that have arisen as a result of the explanatory gaps that manifest within this particular metaphysic. As such, I devote the remainder of this introduction to an articulation of the problems that a complete constitutive panpsychist theory of consciousness must address. I begin with an explanation of what precisely constitutive panpsychism is (and what it is not). From here, I move to contrast a constitutive form of panpsychism alongside other varieties of panpsychism, before arguing that a fully articulated constitutive Russellian micropsychism *probably* offers our best hope of achieving an adequately explanatorily powerful panpsychist theory of consciousness. Finally, I articulate what I take to be the fundamental problems facing constitutive Russellian micropsychism, before hinting at a potential overarching problem that, upon being solved, may lead us to a theory capable of circumventing all explanatory issues for constitutive micropsychism. With this established, I end this introductory section with a (very) brief reiteration of how this thesis will be structured to address this problem, and in so doing hope to lay the foundation for a narrative with the potential to explain precisely how micro-phenomenal properties may be reconciled with our understanding of both ourselves and the natural world.

## 1.1 Constitutive Panpsychism and its place within contemporary philosophy of mind

We each appear to exist as a subject of experience occupying a singular, subjective perspective that is (seemingly) entirely our own and is (seemingly) non-structural in nature. There is ‘something it is like’ (Nagel 1974, p. 442) for each of us to occupy this state, and it is the fact that there *is* ‘something it is like’ to *be* a subject of experience that typifies our existence most indubitably. Further, we seemingly exist as part of a natural world replete with a variety of natural laws that, as far as we can reasonably hope to know, are exclusively causally derivative from the structural properties of the physical universe<sup>13</sup>. In this sense, our perceptual faculties only extend so far as to capture those phenomena which contain, or are constituted upon, structural properties that are physically and spatially extended, and subsequently, as we only have evidence for the existence of such physical phenomena, it might appear only reasonable to infer that all factually existent phenomena *are*, or are at least constituted by, entities that are physically structured in nature<sup>14</sup>. From these two statements pertaining to what we are and what can be said to exist, we can infer two fundamental propositions that, at least *prima facie*, seem to be juxtaposed:

---

<sup>13</sup> Here I am referencing the notion of physical causal closure, which is the physicalist metaphysical claim that every effect is reducible to a physical cause. To deny the causal closure of the physical, we would need an explanation for how a non-spatially extended, non-physical cause could ever produce a spatially extended, physical effect.

<sup>14</sup> For the moment, I note that my interest does not extend to abstract entities such as concepts or numbers (although, as I shall elucidate in chapter 6, we might very well establish a means of reconciling abstractions with the experiential proposition). I also note that an idealist such as Berkeley (1710) would maintain that all properties are mentally constituted and therefore would deny this. For my purposes in this thesis I simply note, in line with Russell (1912), that idealism seems to lead to epistemic stagnation by virtue of denying the existence of the mind-independent objects that act as the ‘epistemic friction’ (McDowell 1996, p. 66) required to justify and ground the content of our thoughts. I do not deem it necessary to articulate the issues associated with adopting a metaphysical idealism in any more depth here (see Moore 1903 and Russell 1912 for convincing arguments against the possibility of all properties being mentally constituted). However, I do note that the fact that contemporary philosophy of mind has moved to supplant metaphysical idealism with more explanatorily robust metaphysics is good reason to reject Dietrich’s (2011) claim that ‘there is no progress in philosophy’ (p. 329).



*Experiential Proposition:* There is ‘something it is like’ to undergo experiential states, and this phenomenal character appears to be non-structural in nature<sup>15</sup>.

*Physical Proposition:* All factually existent phenomena must be constituted by structural, physical properties that account for their spatial extension.

The problem is that these two propositions seem to be dichotomous, yet we cannot easily deny either one without resigning ourselves to a counterintuitive worldview. If we deny the former, we resign ourselves to eliminativism (see Dennett 1991), or a form of the phenomenal concept strategy (see Loar 1990 and Papineau 2002), and in so doing are forced to counter-intuitively deny our existence as non-illusory phenomenally conscious subjects of experience. Alternatively, if we deny the latter, we resign ourselves to a substance or property dualism predicated upon the dubious contention that certain (perhaps phenomenal) properties may exist devoid of spatial extension, but in so doing must inevitably account for how such properties causally interact with the spatially extended physical world. As such, whilst it seems patently true that each of us *does* exist as a phenomenally conscious subject of experience<sup>16</sup>, and each of us *does* exist within a universe in which factually existent phenomena are constituted by physical structures, we face a difficulty upon attempting to reconcile such phenomenal properties with structural, physical properties that, according to

---

<sup>15</sup> In the sense that a description of the nature of phenomenality would not be exhausted by a description of purely narrowly physical, structural properties, or a description of the relation between such structural properties (regardless of the complexity of these relations).

<sup>16</sup> Whilst many would deny this (for examples see Dennett 1991, Loar 1990 and Papineau 2002), for my purposes I do not deem it necessary to offer full articulations of these attempts to maintain the physicalist metaphysic by reducing phenomenal consciousness to an illusion or conceptual misunderstanding. Instead, in line with my aim to advance an explanatorily robust constitutive panpsychist theory, I take the reality of our conscious experiences as a postulate and thereby take phenomenal consciousness to be an actually existent phenomenon that exemplifies a fundamental experiential truth on par with physical truths.

most contemporary iterations of mainstream physicalism (see Dennett 1991), contain a *narrowly* physical, strictly non-phenomenal ontology.

The most standard, but perhaps also most tenuous, way out of this dilemma is to rely upon some form of emergence to explain how narrowly physical, extrinsic properties may constitute intrinsic, phenomenal properties. In this case, we would need an account of emergence that could explain how the ontologically irreducible phenomenal property X *strongly* emerges<sup>17</sup> from the non-phenomenal ontology of physical property Y. However, to explain precisely how the novel ontology of phenomenal properties arose from non-phenomenal properties, we would first need to grant that strong emergence *is* possible, and this is something that is largely disputed. For it seems extraordinarily difficult, if not impossible, to even begin to conceive of a way in which the purely non-experiential, structural properties of matter (regardless of complexity) should ever give rise to the intrinsic experiential property of there being ‘something it is like’ to read this piece, just as it seems equally impossible to conceive of a way in which structural, physical properties might emerge from non-physical, experiential properties. As a resolution to this difficulty, the panpsychist might be read as taking the Humean (1740) notion that inconceivability entails metaphysical impossibility<sup>18</sup> seriously and inferring that as we struggle to conceive of strong

---

<sup>17</sup> *Strong emergence* is best described as the non-deducible, ontologically novel emergence of a new property that is not ontologically reducible to, deducible from, or determined by, the properties of the lower-level domain from which it emerged. In this case, an example of strong emergence would be the novel emergence of phenomenological properties from non-phenomenological, *narrowly* physical properties because such novel properties are ontologically irreducible to the wholly non-phenomenal properties from which they ostensibly emerged. As shall be explored in section 1.1, a strong emergence of this kind is widely held to be impossible, or, at the very least, seriously explanatorily lacking. Conversely, *weak emergence* describes the emergence of a new property that is intelligibly reducible to, determined by, or deducible from the interactions of the properties at the lower-level domain from which it emerged, and for this reason weak emergence avoids the explanatory issues associated with its counterpart.

<sup>18</sup> There is some dispute amongst Hume scholars regarding whether or not Hume endorses an absolute metaphysical impossibility as a result of inconceivability (see Lightner 1997 for an interesting synthesis of the debates in this area). This aside, however, many do cite Hume’s infamous contention ‘that whatever the mind clearly conceives includes the idea of possible existence, or in other words, that nothing we imagine is absolutely impossible’ (Hume 1740, p. 32), and his notion that ‘we can form no idea of a mountain without a

emergence, we ought to advance an alternative solution to the problem of consciousness<sup>19</sup>.

Therefore, when framed in this way, the central panpsychist commitment is as an attempt to advance a solution to the mind-body problem that does not rely upon the possibility that one day strong emergence will bridge the ‘magic passage across the experiential/non-experiential divide’ (Strawson 2008, p. 70).

With this in mind, certain proponents of contemporary panpsychism (see Strawson 2006, Goff 2017, Chalmers 2016) have suggested that the way out of the experiential/non-experiential dichotomy is to expand our concept of the physical<sup>20</sup> so as to accommodate for the possibility that physical properties are both extrinsic, structural qualities *and* intrinsic, phenomenal qualities. In this model, panpsychism might be framed as a form of Russellian Monism, which, as Russell (1948) elucidates, takes the physical proposition that all we can indubitably know is the spatial properties of the physical world seriously, whilst also hinting at a potential means to posit experiential truths into our understanding of the physical:

*The physical world is only known as regards certain abstract facts about its space-time structure – features, which, because of their abstractness, do not suffice to show whether the physical world is, or is not, different in intrinsic character from the world of mind. (p. 240)*

---

valley, and therefore regard it as impossible’ (p. 32) to substantiate the claim that Hume *was* in fact arguing for the claim that inconceivability entails absolute metaphysical impossibility.

<sup>19</sup> I note that this is just one way to motivate Panpsychism, and, in fact, this reading might not provide the strongest foundation for the panpsychist worldview. Indeed, I hold that the explanatory burden that underpins the hard problem might itself simply be enough to motivate an alternative solution to the problem of consciousness which does not require us to either endorse a narrowly physicalist worldview, or embrace the associated explanatory difficulties that occur within a worldview of this kind.

<sup>20</sup> The Panpsychist Russellian Monist is therefore arguing for what Chalmers (2016) terms a *broad* notion of the physical, which is capable of expanding our concept of the physical so as to include more than the *narrowly* physical properties that are exclusively structural whilst maintaining coherence with the proposition that all factually existent phenomena are physically constituted. Similarly, Strawsonian (2006) panpsychism also appeals to extend the reach of physicalism so as to accommodate for the possibility that experiential properties are intrinsic aspects of the physical - in the sense that physical properties just *are* phenomenal properties.

Therefore, a Panpsychist Russellian Monism of this kind offers a neat means of reconciling the experiential and physical propositions by establishing that our understanding of the physical is necessarily incomplete by virtue of being derived purely from an analysis of extrinsic spatial properties, and suggesting that we plug this epistemic hole, whilst simultaneously solving the mind-body problem, by imbuing the physical with a phenomenal quality that accounts for its 'intrinsic character'. However, the extent to which this intrinsic character manifests, and the way it does so, varies based upon the form of Russellian Monism in question. Goff (2017) articulates the types of Russellian Monism as follows:

- *'Panpsychist versus Panprotopsychist - Panspsychist Russellian Monists believe that the deep nature of the physical world is experiential. The typical view is that basic material entities have a very simple experiential nature, from which the complex experience of humans and animals is somewhat derived. Panprotopsychist Russellian Monists, in contrast, hold that the deep nature of the physical is not itself experiential but somehow intrinsically suited for realizing, or bringing about, experience.'*
- *Constitutive versus Emergentist – Constitutive Russellian monists believe that human and animal experience is grounded in, or constituted of, the deep nature of the physical. Emergentist Russellian Monism in contrast holds that human and animal experience is causally brought about and sustained by the deep nature of the physical.*
- *Smallist versus Priority Monism – Smallists believe that all fundamental facts are facts concerning micro-level entities and properties; all things exist and are the way they are in virtue of how micro-level entities are. Priority monists in contrast hold*

*that the universe is the one and only fundamental entity; all other things exist and are the way they are in virtue of how the universe is'. (p. 19)*

For my purposes in this thesis, I will be arguing, in line with Nagel (1979), Seager (2006), Chalmers (2016) and Strawson (2006), in favour of the most standard and widely acknowledged form of Russellian Constitutive Panpsychism: A Constitutive ‘Smallest’ Panpsychist Russellian Monism (or a Russellian Constitutive Micropsychism). In this respect, I am arguing for the metaphysical commitment that the universe is fundamentally extrinsically micro-physically structured whilst being intrinsically microexperiential at a basic level in a manner that ‘involves the instantiation of microphenomenal properties: properties characterizing what it is like to be microphysical entities’ (Chalmers 2016, p. 26), and maintain that facts about macro-level experiential entities are constituted upon these more fundamental and rudimentary micro-level experiential facts. I take it as a given that a Russellian Constitutive Panpsychism of this kind is more resilient than a Non-Russellian Constitutive Panpsychism, for a Panpsychist Russellian Monism is able to account for the causal role of phenomenal qualities in a manner that does not violate our understanding of the causal closure of the physical, by virtue of the phenomenal ‘intrinsic nature’ of the physical being a fundamental part of the causally closed physical world and thereby playing a certain physical role that is manifest in the extrinsic, structural properties of the physical<sup>21</sup>. Conversely, the Non-Russellian Constitutive Panpsychist would need to imbue phenomenal properties with some additional causal function which would be capable of explaining their

---

<sup>21</sup> Chalmers (2018) and Goff (2017) also note this benefit of embracing Russellian Constitutive Panpsychism and maintain that because the microexperiential is grounded in the physical, ‘mind is part of the causally closed physical world’ (Goff 2017, p. 18). Chalmers (2018) describes this by noting that the intrinsic microexperiential nature of the physical may play certain physical roles, such that a ‘microexperiential property [may] realize mass by playing that role, thereby serving as the “intrinsic nature” of mass’ (p. 12). This point pertaining to the causal nature of phenomenal properties shall be greatly expanded upon within chapters 2 and 3.

causal role, but in so doing would face the inevitable difficulty of explaining how this additional causal function may remain consistent with our understanding of physical causal closure. As a result of these explanatory benefits, my focus within this thesis shall rest exclusively upon articulating a Constitutive Russellian Micropsychist theory of consciousness (henceforth referenced as Constitutive Micropsychism). Therefore, in the subsequent section, I contrast a panpsychism of this type alongside the aforementioned variances on Russellian Monism, before arguing that the contemporary movement to supplant Constitutive Micropsychism with Constitutive Cosmopsychism (a la Wager & Nagasawa 2016, Goff 2017 and Shani 2015) is injudicious and, ultimately, concluding that Constitutive Micropsychism offers the potential for the most explanatorily complete means of achieving a panpsychist theory of consciousness.

## **1.2 Why constitutive Russellian micropsychism is (potentially) the most explanatorily complete form of panpsychist Russellian monism**

According to the mainstream understanding of matter advanced by the natural sciences, the fundamental building blocks of our reality are microphysical entities that contain fundamental physical properties. The problems facing the (narrow) physicalist outlook often adopted by the natural sciences are twofold, however: 1) we do not seem to be capable of employing purely physical, non-experiential properties as a means with which to explain our existence as subjects of experience replete with subjective phenomenal characters of ‘what it is like’, and 2) we cannot necessarily rely on our perceptual faculties to delineate the intrinsic nature of the physical (or indeed microphysical) in such a way that would allow us to indubitably demarcate the intrinsic nature of bottom-level microphysical entities. Russellian micropsychism employs this gap in our understanding of the microphysical, and, motivated by the perennial problem of consciousness, endeavours to posit phenomenal properties as

ontologically irreducible intrinsic properties of bottom-level microphysical entities. In so doing, the Russellian micropsychist constructs a worldview that solves the hard problem of consciousness by positioning causally efficacious, *experiential*, bottom-level, microphysical entities as the fundamental constituents of our reality and our conscious experience. In constructing this worldview, however, we inevitably face the troublesome problem of explaining precisely how micro-level experiential entities may simultaneously occupy a variety of individual instances of subjective experience whilst also constituting the unified, macro-level experience we associate with human subjects (this is one way to articulate the ‘combination problem’<sup>22</sup>).

Consequently, to avoid a confrontation with this combination problem several variations on ‘bottom-up’ Russellian micropsychism have been championed. These variances either take issue with the micropsychist commitment to the experiential nature of bottom-level microphysical entities (Panprotopsychism), deny that our macro-level experiential properties are grounded in the experiential properties of the microphysical (Non-Constitutive Panpsychism), or dismiss the notion that the experiential parts of the universe are more ontologically fundamental than the experiential universe itself (Priority Monist Cosmopsychism). As we shall see, however, whilst these variances certainly do seem more resistant to the combination problem, adopting any of these forms of panpsychist Russellian monism necessitates a confrontation with deeper problems that are themselves perhaps more troublesome than the problem they were formulated to avoid. I begin with an articulation of the problems associated with Panprotopsychism.

---

<sup>22</sup> See the subsequent section for a substantially expanded articulation of this problem.

### 1.2.1 Panprotopsychism

Russellian Panprotopsychism is the metaphysical worldview that bottom-level microphysical entities are proto-phenomenal in such a way that, when appropriately combined, the proto-phenomenal properties of the physical may act as the precursor to the macro-level experience we hold within our subjective phenomenal characters of ‘what it is like’. On this view, the world is ultimately constituted of microphysical entities that are proto-phenomenal in nature, such that at the bottom level the universe holds the potential to produce phenomenological experience but is not in itself experiential. So, unlike Russellian micropsychism, which posits that the microphysical universe is microexperiential in such a way that we might say there is ‘something it is like’ to be a quark (if quarks turn out to be fundamental), Russellian Panprotopsychism dispenses of the notion that quarks are micro subjects of experience and instead maintains that 1) proto-experiential bottom-level microphysical entities are non-experiential precursors to phenomenal consciousness and 2) macro subjects of experience arise out of the right conglomeration of proto-experiential microphysical entities. As such, this formation does not posit micro-experiential subjects in such a way that would require for distinct experiential subjects to simultaneously constitute and combine into a singular macro-level subject, and therefore arguably does not fall foul to the combination problem facing constitutive micropsychism.

The problem with this, however, is that in asserting that microphysical entities are *not* implicitly experiential, panprotopsychism must adopt the explanatory burden of accounting for precisely how the *non-experiential* properties of the microphysical may constitute the macro-level *experiential* properties we hold. Indeed, in this respect, and as certain philosophers have highlighted (see Stoljar 2010), the problems facing Russellian



Panprotopsychism appear largely indistinct from the problems facing narrow physicalism, as both begin with non-experiential microphysical entities and endeavour to traverse the ‘magic passage across the experiential/non-experiential divide’ (Strawson 2008, p. 70).

Consequently, it seems that positing that the microphysical is proto-phenomenal, but (somehow) holds the potential to produce experiential properties proper, is no different from championing the physicalist line that the microphysical is non-experiential, but (somehow) holds the potential to produce experience. Therefore, *prima facie* it appears that constructing a metaphysical worldview based upon the notion that microphysical entities are proto-phenomenal leaves us no better equipped to solve the problem of consciousness than if we were to embrace narrow physicalism.

To avoid this problem those sympathetic to panprotopsychism have argued that we must sharpen our definition of proto-phenomenal properties to highlight the implicit link between proto-experience and experience proper in such a way that we could a priori extrapolate macro-experiential truths from proto-phenomenal truths (see Goff 2015, Chalmers 2013).

Chalmers (2013) endeavours to distance panprotopsychism from the problems facing narrow physicalism by prefacing the panprotopsychist worldview as follows: ‘(i) proto-phenomenal properties are distinct from structural properties, and (ii) there is an a priori entailment from truths about proto-phenomenal properties (perhaps along with structural properties) to truths about the phenomenal properties that they constitute’ (p. 16).

One of the most promising means of making sense of these conditions is found within Coleman’s (2016) articulation of ‘Panqualityism’, which posits that the world is ultimately ‘constituted of quality instances’ that exist as ‘unexperienced qualia – properties just like the qualities we experience, but without anyone experiencing them’ (Coleman 2016, p. 249). On

this view, upon seeing the blackness of these letters as contrasted against the white backdrop, we are undergoing an experience of the phenomenal qualities involved in seeing (or perhaps in this case *not* seeing) colour, and it is exclusively these qualities that the panqualityist takes to be the protophenomenal properties of the microphysical. So, whilst panpsychism takes the subjective experience itself to be ontologically fundamental, the panqualityist attempts to posit the qualities of experience as fundamental, and endeavours to extrapolate a reductive account for how experience arose out of unexperienced qualities.

Such an account faces two arguably insurmountable issues. First, it is not at all clear that it is possible to make sense of the concept of unexperienced qualia, for it seems extraordinarily difficult to account for how a quality of experience can exist without an experiencer<sup>23</sup>. Second, Coleman (2016) seems to persistently infer that panqualityism faces another version of the combination problem, which must be addressed to account for how ubiquitous unexperienced phenomenal qualities may combine to form the macrosubjectivity we experience (p. 249-251). I take issue with this, for it seems that this is fundamentally not a ‘combination’ problem; this is a generation problem. Panprotopsychism, or Panqualityism, both begin by explicitly *not* positing microphenomenality as an ontologically fundamental property of the physical, and therefore must account for precisely how the non-experiential microphysical generates subjective conscious experience. To infer that this is a combination problem is to infer that phenomenal experiences ‘Y’ are ontologically reducible to unexperienced qualia ‘X’, and it seems extraordinarily difficult, if not impossible, to see how this may be the case without relying on the sort of strong emergence that Russellian accounts of Panpsychism were designed to avoid in the first place. In this respect, I argue that despite

---

<sup>23</sup> I shall substantially expand on the nature of this problem in the subsequent chapter.

Chalmers' and Coleman's efforts, protophenomenal properties remain largely ambiguous because it remains unclear precisely what is being added to the standard physicalist account of bottom-level properties to imbue reality with the *potential* to produce conscious experience. If panprotopsychism posits any form of minimal consciousness at the bottom-level at all, then it invariably violates its central tenet by rendering itself indistinct from a more standard form of micropsychism, however, if it persists in maintaining that ambiguous unexperienced phenomenal properties exist at the bottom-level, then arguably Coleman's (2016) panqualityism is not impervious to precisely the problems facing narrow physicalism, because if the bottom-level is *not* phenomenal, it is not at all clear how phenomenality might occur without reinviting the issues entailed by strong emergentism. We might even say that narrow physicalism, replete with all its explanatory issues, seems preferable to panqualityism in this respect, for at least the former is explicit on the intrinsic nature of the physical, whereas the latter seems to adopt all the generation problems of the former whilst also rendering the nature of the physical entirely unintelligible<sup>24</sup>. In this respect, the strength of Constitutive Micropsychism appears immediately transparent, for this metaphysic allows us to embrace the advantages of worldviews like physicalism which are explicit on the ontology of their constituent parts, whilst also allowing us to avoid a confrontation with the problem of generating experiential properties from non-experiential properties. With the issues facing panprotopsychism articulated, I now move to address some of the problems associated with a Non-Constitutive Panpsychist Russellian Monism.

---

<sup>24</sup> To avoid this charge, the panqualityist would need to advance a convincing argument to address precisely how unexperienced qualia can be coherently conceptualized. However, as I articulate within the next chapter, such an argument is currently elusive.

## 1.2.2 Non-constitutive Panpsychism

As distinct from constitutive forms of panpsychism, which take it as a given that the intrinsic nature of macro-level conscious experience is constituted by the intrinsic experiential nature of microphysical entities, a Non-Constitutive Russellian Panpsychism holds that facts about macro-level consciousness are *not* constituted, or grounded in, facts about microexperience. In this sense, whilst Constitutive Panpsychism posits that our macro-level conscious experience is non-fundamental and is therefore constituted by more fundamental, experiential, bottom-level microphysical entities that (somehow) combine into a macro-level conscious subject, Non-Constitutive Panpsychism maintains that human and animal consciousness is itself fundamental in such a way that the experiential nature of the quarks or photons that constitute our physical brain does *not* constitute the macroexperiential nature of conscious subjects. In this respect, the crucial difference between these competing worldviews is that the former argues in favour of the ontological reducibility of macro-experience to micro-experience, whereas the latter argues that human and animal consciousness is entirely irreducible to micro-level experience.

Subsequently, in this latter case, the non-constitutive panpsychist seems reliant upon some form of strong emergence to explain precisely how the ontologically novel properties associated with macro-level conscious experience emerge from micro-level experiential properties. Advocates of non-constitutive forms of panpsychism (see Brüntrup 2016) maintain, however, that this does not leave us entrenched in the same problem facing narrow physicalism, for we can distinguish the emergence entailed within this worldview from the ‘superstrong emergence’ entailed by the ‘emergence of a phenomenal mind from a [narrowly physicalist] world which is merely spatial extension or a framework of causal-functional

interactions' (Brüntrup 2016, p. 69). In this respect, non-constitutive panpsychists tend to argue that as bottom-level microphysical entities still contain phenomenal properties, we do not necessarily require the emergence of entirely novel properties to explain the emergence of macroexperience from microexperience, and we therefore avoid the unintelligibility of a physicalist worldview in which novel experiential properties emerge from a lower-level domain that is entirely non-experiential. However, whilst this may be true, the fact that macroexperience is *not* reducible to microexperience still leaves the non-constitutive panpsychist in need of an explanatorily consistent account of precisely how macroexperience emerged.

Here, the theories that have seen the most philosophical success are the Synchronic Non-Constitutive Emergence endorsed by Brüntrup (2016) and the Diachronic 'Combinatorial Infusion' championed by Seager (2010). I shall begin by elucidating the former. On this view, microexperiential properties causally necessitate the existence of macroexperience in a manner that allows for microexperiential subjects to simultaneously cause and exist independently of macrosubjectivity, such that once microexperiential properties are amalgamated, or arranged in a specific way, these properties will cause the emergence of macrosubjectivity, which exhibits a distinct ontology and exists alongside microexperiential properties. On the opposing view, microexperiential subjects fuse together in such a way that simultaneously causes microsubjectivity to cease to exist and causes the emergence of macrosubjectivity. In this respect, the difference between these two versions of Emergentist Panpsychism lies in whether microsubjectivity continues to maintain an independent existence, or ceases to exist, upon the emergence of macrosubjectivity.

Of these two accounts, I propose that the latter is the most robust, for if we embrace the former and thereby remain open to the possibility that microsubjectivity causes macrosubjectivity whilst maintaining an independent existence, it seems we need an adequate account of precisely how microsubjectivity can co-exist with macrosubjectivity without facing a generation problem. This is because it is difficult to conceive of how a myriad of instances of subjective experience can cause an entirely new instance of unified subjective experience whilst remaining individual instances of subjectivity. It seems this new macrosubjectivity must emerge from *something* (i.e., the microsubjectivity itself), but this strays somewhat close to the magic of ‘something from nothing’ if the instances of microsubjectivity remain entirely unchanged. To avoid this, this form of Emergentist Panpsychism is usually framed as an Emergentist Panprotopsychism, in which the microphysical simply holds proto-phenomenal properties that facilitate macrosubjectivity. As a result, the advocate of Synchronic Non-Constitutive Emergence does perhaps avoid the problem of how to account for the co-existence of micro and macro subjectivity, however, in so doing seems to face precisely the generation problem that afflicts all instances of panprotopsychism; for, if microphysical entities are proto-experiential, it seems we must account for how the non-experiential microphysical generated experience, and as a result must resign ourselves to the ‘superstrong’ emergence that the non-constitutive panpsychist was attempting to avoid.

Combinatorial Infusionism arguably presents the stronger form of Emergentist Panpsychism, then. On this worldview, the microphysical is feasibly entirely subjective in nature at time 1 but, upon fusing with other instances of microsubjectivity at time 2, a novel, unified macrosubjectivity occurs and thereby so too does the cessation of microsubjectivity. Seager (2010) describes this view as follows:

- ‘1. The mental character of the combined or aggregative mental state stems from the mental characteristics of the constituents,*
- 2. The combined or aggregative mental state is a novel state which in some way ‘absorbs’ or supersedes the mental states of the constituents, but*
- 3. There is no radical or ontological emergence of the aggregative mental state; rather there is an intelligible relation which holds between the mental components and the resulting aggregative state.’ (p. 180)*

Seager (2010) employs the example of a black hole in order to explain this theory and maintains that just as the gravitational collapse of disparate physical entities renders the simultaneous cessation of such entities and the fusion of a novel black hole with new properties, the brain exists as a novel macro-level ‘large simple’ (p. 180) that manifests new properties that are caused by the fusion of lower-level physical entities, but remains irreducible to the properties found at the lower-level domain and no longer exemplifies any trace of such lower-level physical entities.

The problem with this is twofold, however. Firstly, it seems difficult for Seager to coherently defend the claim that ‘that there is no radical or ontological emergence’ whilst simultaneously maintaining that emergent macro-level ‘large simples’ exemplify entirely new entities with irreducible ‘novel states’ (p. 180), for if the macro-level ‘large simple’ is entirely novel in nature in such a manner that it manifests a new property that is entirely irreducible to the properties of its lower-level constituents, then it seems we are necessarily discussing radical emergence. This is important because radical emergence is something that Seager (2010) is very keen to avoid - as Seager himself notes, if his fusionism does rely upon radical

or ontological emergence, it adopts all of the more standard issues associated with a ‘strong’ emergence of this kind and therefore must explain precisely how the entirely novel property ‘Y’ arose from nothing other than a quantity of complexly interacting lower-level properties ‘X’, all of which are ontologically distinct from Y (see p. 177-80) <sup>25</sup>.

Alternatively, if Seager persists in denying that these novel properties are truly radically or ontologically emergent, then we are arguably discussing a form of constitutive micropsychism in which an entity (weakly) emerges that is reductively accounted for by the relations between the properties of its constituent parts. If this were the case, Seager (2010) would be able to intelligibly make sense of the claim that there is ‘no radical emergence here, but there is nonetheless the creation of a new entity’ (p. 182), however in so doing must clearly forego not only the notion that this new entity exists as an ontologically novel ‘large simple’ in which entirely non-reducible or deducible novel properties manifest, but also his commitment to the central tenet of Non-Constitutive Panpsychism<sup>26</sup>. Whilst such a move would render Seager’s fusionism wholly consistent with the aims of this thesis, the problem with this is that, as Seager himself points out, ‘since combinatorial infusionism is a diachronic relation it is hard to see how it could be constitutive’ (Seager 2016, p. 198). In this vein, I maintain that whilst we might be able to conceive of other accounts of fusionism that are consistent with constitutive micropsychism (see chapter 5), the only way to make sense of Seager’s account of combinatorial infusionism is to posit it as a type of diachronic, non-

---

<sup>25</sup> For my purposes, a more detailed examination of the problems associated with ontological, or ‘strong’, emergence is unnecessary. All such problems can be reduced to the difficulty of explaining precisely how entirely ontologically irreducible, novel properties can emerge from a lower-level domain that does not ontologically necessitate or explain the emergence of such properties. For detailed examinations of this problem see Bitbol (2007) or Heard (2005).

<sup>26</sup> With this said, I reiterate that Seager *could* conceivably avoid the issues afflicting his account by attempting to explicitly posit Combinatorial Infusion as a type of Constitutive Panpsychism (see Seager 2016 for an articulation of precisely this strategy).



constitutive, strongly emergent panpsychism, for it seems the novel ‘large simple’ cannot be constituted from the properties of lower-level entities whilst simultaneously only existing by virtue of the fusion of such entities in a manner that 1) necessarily causes the cessation of the lower-level entities and thereby their properties, and 2) causes the emergence of a novel property.

The second problem with Seager’s account, as Papinaeu (2001) points out, is if new entities with novel properties truly do strongly emerge from lower-level constituents, we should expect to witness a ‘display of some manifestation of their presence’ (Papinaeu 2001, p. 31), for it is entirely reasonable to expect that, if such strongly emergent properties exist, they would manifest new causal forces that are explicitly distinct from the forces manifested by the lower-level entities from which they emerged. Subsequently, it is reasonable to expect that they would imbue reality with causally efficacious novel properties that manifest as forces that interact with physical laws in a manner that differentiates them from the lower-level domain. However, as the scientific method has consistently revealed that the same physical laws apply ubiquitously to both macro and micro level phenomenon, it seems reasonable to infer that there are no truly strongly emergent novel properties. In this respect, both Papinaeu (2001) and Goff (2017) take the fact that as ‘cellular biology and neurophysiology have never revealed any sign of emergent forces [we have] a strong inductive case to think that there are no such forces’ (Goff 2017, p. 245), and I would agree. Therefore, the central problem facing the non-constitutive panpsychist is one of reconciling the claim that strongly emergent properties exist with the lack of empirical evidence for the existence of such emergent properties, for without this non-constitutive panpsychism distances us from the widely accepted notion of scientific reductivism without presenting an empirically robust alternative.

In light of these difficulties, the constitutive panpsychist (see Goff 2017) argues that we ought to exclusively endorse either a ‘Smallest’ Constitutive Micropsychism or a Priority Monist Cosmopsychism, as both allow us to explain the (weak) emergence of the mind in a manner that avoids the empirical issues facing the non-constitutive panpsychist. Indeed, the empirical robustness of constitutive panpsychism is largely conceded by non-constitutive panpsychists, as ardent non-constitutive panpsychists such as Brüntrup (2016) elucidate, the beauty of Constitutive Panpsychism, as opposed to its non-constitutive counterpart, lies in its simplicity:

*The answer [to the question of how the mind emerged] is straightforward for the constitutive panpsychist: The individual cells have intrinsic natures that are mental or at least analogous to mentality (that is, proto-mental)<sup>27</sup>. It is the composition of these intrinsic natures that explains the emergence of phenomenal minds. The emergence is thus neither brute nor inexplicable. The composition of the cells alone accounts for the weak emergence of higher-level structures and higher-level mentality. There is logical synchronic supervenience between the lower and the higher levels. A perfect copy of all of the cells, including their intrinsic natures, will necessitate higher-level structure and higher-level phenomenal properties. (p .59)*

Thus, whilst non-constitutive forms of panpsychism are charged with addressing the inevitable ontological and empirical issues surrounding the strong emergence of novel properties, the constitutive panpsychist offers a relatively straightforward solution to the

---

<sup>27</sup> I note that in my conception these intrinsic natures are strictly phenomenal, and therefore I avoid the ambiguities surrounding accounts such as Brüntrup’s which introduce the notion of the ‘mental’.

problem of consciousness that leaves scientific reductivism intact via simply adding ‘nonobservable intrinsic natures to the scientific image’ (Brüntrup 2016, p. 59). As a result, and in line with my aim of advancing an argument in favour of construing constitutive micropsychism as the most explanatorily complete form of Russellian panpsychism, I devote the remainder of this section to an elucidation of the problems associated with the one remaining viable alternative to Russellian micropsychism: Priority Cosmopsychism.

### 1.2.3 Priority Cosmopsychism

Unlike the aforementioned articulations of Russellian panpsychism, which inevitably adopt the issues entailed by relying upon some form of strong emergence to explain the phenomenon of conscious experience, priority cosmopsychism is a type of constitutive panpsychism and therefore avoids said issues by maintaining that human conscious experience is constituted by the intrinsic experiential nature of the physical. However, as distinct from a constitutive micropsychism, which is grounded in the concept of atomism and posits microphysical bottom-level particles as the ultimate constituents of our reality, this metaphysic adopts a holistic outlook<sup>28</sup> in an attempt to transition away from the ‘bottom up’ approach entailed by micropsychism by combining panpsychism with holism to elucidate a holistic model for a conscious universe in which our subjective conscious experience is grounded in the intrinsic nature of the macro-level ‘cosmic-consciousness’, and not in the intrinsic nature of microphysical entities.

---

<sup>28</sup> In this context, this is the claim that we exist as part of a ‘top down’ universe in which wholes are more fundamental than their parts.

This model draws inspiration from the works of physicists (see Bohm 2005, Peat 1987, Wolfram 2002) and philosophers of science (Smuts 1926) who have attempted to maintain the intuitive notion that there *is* a metaphysical foundation to our reality<sup>29</sup>, i.e., a foundational ‘bottom-level’ in which fundamental entities exist that are no longer ontologically dependent upon other entities, but have hoped to supplant the reductionism entailed by an atomistic worldview, in which the properties of complex structures are reduced to the properties of ontologically fundamental ‘bottom-level’ microphysical entities, by arguing that because we can make no sense of a singular phenomenon (such as elementary microphysical entities) existing outside of the complex universe, we must infer that such particles are simply parts, or aspects, of the complex universe itself, and therefore the complex universe, and not the microphysical sub-atomic particles, acts as the metaphysical foundation, which is indivisible and thereby more fundamental than its parts. Holism, therefore, stands as the antithesis to atomism by circumventing the reductionist view that fundamental microphysical particles constitute macro-level entities and instead advances a model for a top-down, antireductionist worldview in which particles, human brains, and all other physical phenomena derive their existence from the ultimate macro-level, ontologically primary whole: the universe itself. In this respect, the microphysical sub-atomic particles that exist within this page should *not* be thought of as ontologically fundamental and should instead be construed as simply aspects of the page itself, and the page itself should be thought of simply as an aspect of the ontologically fundamental macro-level universe.

---

<sup>29</sup> I refer to this as intuitive because the alternative, ‘metaphysical infinitism’, entails that there is no ontological foundation to reality (and therefore no real means of delineating the ultimate nature of reality), for on this account all of reality is reduced to an indivisible infinite regress of ontologically mysterious ‘gunk’ (Schaffer 2010, p. 36). See Tahko 2014 for a full articulation of the disparities between foundationalism and infinitism, and an articulation of why foundationalism is simply more intuitive than infinitism.

With this as grounding, priority cosmopsychists such as Goff (2017) and Nagasawa & Wager (2016) extrapolate that 1) *if* holism is true and all physical phenomena derive their existence from the universal whole, and 2) *if* Russellian panpsychism is true and the intrinsic nature of the physical is experiential, then the universal whole is necessarily experiential, and therefore the experiential nature of the human brain is simply an aspect, or part, of the experiential cosmos. Goff (2017) maintains that because cosmopsychism allows for human conscious experience to be explained as simply a part, or aspect, of an ontologically primary, macro-level experiential whole, it holds an advantage over its micropsychist counterpart by avoiding the need to explain precisely how many individual instances of micro-subjectivity simultaneously constitute and combine into a singular macro-level subject of experience. However, whilst this may be true, cosmopsychism is not without its own set of explanatory hurdles.

To render said hurdles transparent, let us elucidate the worldview that the priority cosmopsychist champions. It consists of a singular cosmic consciousness replete with its own subjectively unified experience of ‘what it is like’, and a myriad of instances of subject-specific consciousness derived from the underlying ubiquity of the cosmic consciousness. This means that the subjective consciousness of the cosmopsychist is necessarily constituted from the subjective consciousness of the cosmos. Yet, the cosmopsychist is swift to affirm that his subjective experience is something distinct from the subjective experience held by the cosmos and is instead simply a conscious ‘aspect of the whole’ (Goff 2017, p. 235), in such a way that if the cosmos is in one ‘determinate state of consciousness - I will be in another determinate state of consciousness’ (Goff 2017, p. 236). Indeed, this is seemingly something the cosmopsychist *must* maintain, for to attest otherwise resigns the cosmopsychist to the claim that instances of subjective experience are not subjective at all, and therefore strays

dangerously close to a denial that we exist as distinct experiential subjects replete with private phenomenal characters.

As the cosmopsychist cannot easily deny the internal privacy of our subject-specific experience, he must provide an account of how the cosmic consciousness of the universe first split itself into distinct experiential aspects in such a way that would adequately explain how the universe was able to experientially alienate itself from itself, and subsequently must also explain how there is now a multiplicity of experientially distinct subjects that are capable of experiencing the universe as a mind-independent object. Accounting for how the universe performed this ‘first split’ seems to be the very core of the problem facing the priority cosmopsychist, for it seems extraordinarily difficult to conceive of precisely how the singularly experiential conscious cosmos at time 1 (T1) split to create a multitude of new subject-specific conscious aspects at T2 that exist simultaneously as a part of the cosmic consciousness but are also able to observe the cosmos as an object distinct from themselves. As Kastrup (2018) and Coleman (2014) elucidate, the problem is one of explaining precisely how lower-level, experientially ‘relative’ subjects (i.e., subjects with a private subject-specific experience), form within the conscious cosmos. Or, more simply, if the conscious cosmos is in one determinate, subject-specific state of consciousness, how precisely do lower-level subjects acquire their own subject-specific, private point of view? (Coleman 2014, p. 30-37 and Kastrup 2018, p. 131-134). Indeed, without an adequate explanation for precisely how the universe split itself in this manner, the cosmopsychist is arguably left confronting a more damning problem than the more standard subject-summing problem facing micropsychism.

Spencer-Brown (1969) wrote before cosmopsychism entered the contemporary debate but still manages to elucidate this problem neatly by maintaining that, for the world to perceive itself, ‘it must first cut itself up into at least one state that sees, and at least one state that is seen’ (p. 104)’. The problem of accounting for precisely how the universe may cut itself up in this way is what Spencer-Brown termed the ‘original mystery’ (p. 105):

*‘In order for the world to see itself...it must first cut itself up. In this severed and mutilated condition, whatever it [the world] sees is only partially itself. We may take it that the world undoubtedly is itself (i.e. is indistinct from itself), but, in any attempt to see itself as an object, it must, equally undoubtedly, act so as to make itself distinct from, and therefore false to, itself. In this condition it always partially eludes itself. It seems hard to find an acceptable answer to the question of how or why the world discovers an ability to see itself, and appears to suffer the process.’* (Spencer-Brown, 1969, p. 104-5)

Whilst Spencer-Brown’s aim in this passage was to highlight the explanatory difficulties facing the natural sciences, his sentiments seem to be particularly scathing considering the recent move to champion priority cosmopsychism as an attractive option within the field of the metaphysics of the mind. Especially as this seems to be the first robust iteration of what is now termed the ‘decomposition’ (Chalmers 2016, p. 196), or ‘decombination’ (Kastrup 2018, p. 125), argument against cosmopsychism. This problem of explaining precisely how a singular cosmic consciousness, replete with its own unified and subjective phenomenal character of ‘what it is like’, was able to experientially alienate itself from itself by splitting itself into a myriad of instances of subjective-specific consciousness of the type we experience seems to be just as troublesome, if not more so, than the original articulation of the combination problem charged against micropsychism. Both such problems are reducible to a conceptual issue of imagining precisely how conscious subjects can either constitute or

contain other conscious subjects, and therefore, as Miller (2018) articulates, this decombination problem is at least equivalent to the combination problem, for both are concerned with ‘subjects being proper parts of other subjects’ (p. 137).

Miller (2018) maintains that the root of this decombination problem can be traced to the dichotomy that arises amidst cosmopsychism and our simultaneous belief that there are two defining qualities of a phenomenally conscious subject: 1) we each exemplify a phenomenally unified experiential locus that is entirely subjective in a manner that all of our various experiences are construed as parts, or aspects, of a unified phenomenal field<sup>30</sup>, and 2) this subjective state is entirely private in such a way that experiences and thoughts belong, or are ‘bound’, exclusively to this particular subject ‘S’ in a manner that could not be replicated amidst those subjects ~ S. With this established, there does appear to be a real difficulty involved in accounting for how the singular experience of the ‘subject-whole’, i.e., the experientially unified, subject-specific experience of ‘what it is like’ for the cosmos, may ground phenomenally ‘unified and bounded’ subjects replete with their own phenomenal experiences. This difficulty, as Miller (2018) notes, arises precisely because the cosmos cannot simultaneously be an entirely unified subject of experience, replete with its own subject-specific experiential locus, whilst simultaneously constituting a myriad instances of unified and bounded subject-specific experience of the type we associate with human consciousness, for the bounded subject-specific (or private) experience of the latter seems to necessarily contravene the subject-specific experiential unity of the former<sup>31</sup>.

---

<sup>30</sup> This notion seems to be borrowed from Bayne (2010).

<sup>31</sup> By this Miller (2018) means that because ‘we cannot project ourselves into the perspective of a phenomenally scattered subject...phenomenal unity must be a necessary and constitutive feature of a subject’s consciousness’ (p. 153), and therefore the universe cannot simultaneously exist as a subject of experience whilst constituting a myriad of relative, subject-specific experience, because this seems to necessarily produce a phenomenally scattered subject, and this is something that we simply cannot coherently comprehend.



To address this problem, the cosmopsychist has three conceivable options: 1) maintain, as Kastrup (2018) does, that ‘dissociation in cosmic consciousness is what leads to the formation of relative subjects’ (p. 142), 2) uphold that the universe may well exemplify a privately subjective experiential locus at T1, but, by virtue of submergence, an entirely new multiplicity of subjective experiential loci arose at T2, or 3) maintain, as Goff (2017) does, that ‘what the cosmopsychist takes to be fundamental is not simply a conscious universe, but *a conscious universe which contains other conscious subjects as partial aspects*’ (p. 118).

Of these three solutions, both latter options contain immediate issues that must be addressed. First, those cosmopsychists that champion the notion that subject-specific consciousness is a submergent property of the universe must address precisely the same problems facing the iterations of diachronic or synchronic emergence espoused in the previous section, and therefore seem to necessarily resign themselves to a worldview that is largely indistinct from the more standard forms of non-constitutive panpsychism. Secondly, for Goff to maintain that the universe fundamentally contains ‘other conscious subjects as partial aspects’ is to render cosmopsychism indistinct from the ‘bottom-up’ micropsychism it is presented as an alternative to, for as soon as we move from a singular fundamental conscious subject to a multiplicity of fundamental conscious subjects, we are arguably no longer discussing a coherent form of cosmopsychism by virtue of allowing for our ontological foundation to contain a multiplicity of subject-specific experience<sup>32</sup>. Even if the cosmopsychist resists this charge by emphasising that individual instances of phenomenal experience are simply aspects of the fundamental conscious cosmos (and are thereby *not* instances of ontologically

---

<sup>32</sup> I note here that Shani (2015) also falls foul to this particular issue.

fundamental consciousness in their own right), it seems he must still account for how, if the conscious cosmos exists as a unified and bound subject of experience, it may contain an array of relative (or ‘bounded’) subject-specific aspects, because this would render the consciousness of the cosmos either unbounded (in the sense of non-subject-specific), or *incomplete* (in the sense of not being an instantiation of ‘what it is like’ for the universe taken as a whole). The central point here is that, if cosmopsychism is true and the universe undergoes a subject-specific, private, ‘what it is like’ experience that is complete and ‘whole’ in the sense that it is a unified exemplification of every aspect of the universe, the occurrence of additional experientially ‘bounded’ (I take bounded to be synonymous with isolated/private/relative), subject-specific conscious aspects necessarily implies that the cosmic consciousness is either 1) not complete, in the sense that the experience of the cosmic ‘whole’ is *not* an exemplification of all aspects of the whole (and therefore the gaps in the experience of the cosmic consciousness seem to violate the cosmopsychist commitment to holism by allowing for the possibility that there are aspects to the universe that are *not* grounded in the holistic whole), or 2) the cosmic consciousness is not experientially bounded, in the sense that whilst the experience of the cosmic unified ‘whole’ might be ‘complete’ in a manner that it exemplifies all aspects of the universe whilst maintaining coherence with Goff’s (2017) notion that the cosmic consciousness contains ‘conscious subjects as partial aspects’ (p. 118), in so doing it must be devoid of subject-specific experiential boundaries<sup>33</sup>, because if such boundaries are in place the cosmic consciousness can never truly encompass a unified representation of the ‘whole’ universe. However, if it is the case that Goff’s cosmopsychism leads to unbounded subjects, the cosmic consciousness is arguably *not* an exemplification of subject-specific experience at all (and thereby arguably not an adequate

---

<sup>33</sup> Here I am referencing the subject-specific experiential boundaries that separate a subjective experience for S at T1 as distinct from the experience held at T1 by all subjects ~ S.

exemplification of phenomenal experience)<sup>34</sup>. This seems to be necessarily true, for it is only by virtue of subject-specific boundaries that we can demarcate an experience *for* subject S at T1 as distinct from the experiences held by subjects  $\sim$  S at T1. Without such boundaries, it becomes incredibly difficult to conceive of how we can make sense of the concept of subjects of experience (or indeed the concept of subjectivity), for such concepts seem to necessarily imply that there is a boundary that separates an experience for S as not simultaneously an experience for all subject  $\sim$  S. As such, either the ontologically fundamental cosmic consciousness *does* contain relative, subject-specific aspects replete with their own private phenomenal experiences that are *not* shared with the universe, and therefore these isolated conscious aspects render the experience of the cosmic consciousness incomplete in a way that violates the central holistic tenet of cosmopsychism, or the cosmic consciousness does *not* contain relative, subject-specific aspects, and therefore the cosmopsychist struggles to explain our existence as phenomenally conscious subjects of experience by virtue of endorsing an ontological foundation that does not coherently incorporate subjectivity. This leaves the cosmopsychist reliant upon a solution that advocates both 1) that a singular, subject-specific cosmic consciousness is ontologically fundamental, and 2) that all iterations of relative, subject-specific consciousness are (somehow) grounded in the cosmic consciousness. I now turn to Kastrup's (2018) dissociation thesis.

---

<sup>34</sup> I note that the core of this inference might be made clearer in the subsequent chapter. In short, here I am referencing the difficulty of accounting for how we may make sense of phenomenological experience without the existence of a subject for whom there is 'something it is like' to undergo an experience. In this sense, if the universe is unbounded and therefore not an exemplification of 'subject-specific' consciousness, it becomes difficult to make sense of how cosmopsychism may solve the 'hard' problem of explaining our existence as subjects of experience replete with phenomenal characters of 'what it is like', for such an explanation seems to necessarily require positing phenomenal experience, and thereby subjectivity, at the outset (a la panpsychism), or advancing some form of emergentism. As cosmopsychism is an example of the former, it seems intuitive that advocates of this type of panpsychism would be reluctant to endorse the idea that the cosmic whole does not hold subject-specific phenomenal experiences.

Kastrup (2018) maintains coherence with these conditions by beginning with the claim that ‘there is only [one] cosmic consciousness’ (p. 142) that is ontologically fundamental, and explaining that all relative, subject-specific aspects of this universal consciousness are themselves simply instances of the cosmic consciousness that have become dissociated from the cosmic consciousness which grounds them, in such a way that ‘an illusion of separation arises’ (p. 143) between the identity of subjectively conscious aspects and the conscious cosmos. Kastrup relies heavily on the concept of dissociative identity disorder (DID) in order to ground this thesis and ultimately aims to solve the decombination problem by asserting that just as the sufferer of DID exhibits a multiplicity of distinct centres of consciousness, so too does the cosmic consciousness afflicted by DID exhibit a multitude of centres of conscious experience (in line with the psychiatric literature in this area, Kastrup references such centres of conscious experience as ‘alters’). In this respect, in no uncertain terms, Kastrup argues that ‘dissociation in cosmic consciousness is what leads to the formation of relative subjects’ (p.143), as, according to Kastrup, it is only by virtue of this process of cosmic consciousness alienating itself from itself that we may make sense of the concept of the subject-specific conscious universe splitting itself into a multitude of subject-specific experiential subjects that act as both perceivers of objects and objects of perception.

The problem with this line of argument, however, is that it is not immediately clear precisely how we may employ the concept of DID to explain the occurrence of phenomenally separated subjects of experience that are subjectively and privately bounded in such a way that they are completely phenomenally isolated from the cosmic consciousness. This is important because it is the problem of explaining how new subjects of experience may become phenomenally isolated from the cosmic consciousness that grounds them that typifies the very core of the decombination problem, yet it is an explanation of this kind that eludes

Kastrup, for without relying upon some account of strong emergence it remains inexorably difficult to coherently make sense of the notion that a singular, bounded, subject of experience 'S' constitutes a multitude of subjectively bounded conscious aspects that are phenomenally isolated from S in such a way that they may be coherently and accurately demarcated as phenomenally new subjects that are  $\sim S$ .

Kastrup (2018) realises this, and asserts that, naturally, 'because alters are fully grounded in cosmic consciousness, it is incoherent to say that they become separated from it; only an illusion of separation arises as a particular phenomenal content in the alter's dissociated qualitative field' (p. 142-143). This leaves Kastrup with a further problem, however, for whilst Kastrup must maintain that there is only the illusion of separation between S at T1 and the conscious aspects of S at T2, in so doing Kastrup renders the concept of subjectivity, and indeed our most fundamental intuitions about the nature of conscious experience, entirely illusory.

This is because, for Kastrup's thesis to maintain internal coherence, the boundaries that separate the experiences for the cosmic consciousness S as distinct from the experiences held by the conscious aspects of S can never produce new instances of actual, non-illusory subjectivity, because for the subjectivity of conscious aspects to be non-illusory such aspects would have to be decidedly  $\sim S$ . However, as all conscious aspects are grounded in the subject-specific consciousness of the cosmos 'S' in a way that does not lead to the ontological possibility of conscious aspects arising that are  $\sim S$ , this is something that cosmopsychism simply cannot account for, and therefore if the boundaries that enable us to coherently conceptualise subjective conscious experience are themselves entirely illusory, then (if we embrace this model) our most fundamental intuitions about the subjective nature

of our phenomenal experiences must also necessarily be illusory. This is a problem for Kastrup, for it seems that, in rendering subjectivity illusory, he faces the difficulty that the concept of a conscious cosmos afflicted with DID seems to not only deny the existence of the phenomenon it was purporting to explain by explaining away the issue of how the cosmos' singular phenomenal experience of 'what it is like' grounds a multitude of distinct, non-illusory, subject-specific experiential boundaries (and thereby seems to entirely avoid the fundamental problem of explaining subjectivity), but also violates our deeply held intuitions about the non-illusory nature of subject-specific experience. In this regard, it seems Kastrup (2018) commits himself to a line of argument that is not overtly distinct from Dennett's (1991) attempt to posit phenomenal experience as nothing beyond an illusion<sup>35</sup>, and, in so doing, seems to invite the charge that explaining away subjectivity does not correspond to an effective explanation for subjectivity. As a result of such difficulties, if we wish to take our most fundamental intuitions about the nature of subjective experience seriously, it seems far more parsimonious to simply attest that our ontological foundation does *not* contain a singular subject-specific conscious experience, but instead contains a multiplicity of subjective-specific experience (as is entailed by constitutive micropsychism), for such an account can entirely avoid the decombination problem by easily maintaining that the universe has always been split in such a manner that allows us to account for the occurrence of both perceivers of objects and objects of perception.

Whilst these issues alone seem to offer coherent reasons to resist the recent move to supplant micropsychism with cosmopsychism, I argue that even if we were to forego these explanatory

---

<sup>35</sup> I note here that Kastrup's thesis is arguably less demanding than Dennett's, for whilst the latter argues that all phenomenal experience is illusory, the former argues that only our perception that we exist as distinct, subject-specific experiential loci is illusory. I do maintain, however, that, as both require us to dispense of our most fundamental intuitions about the nature of consciousness on the grounds that they are simply illusory, both may coherently be charged with explaining away the very problem they were purporting to solve.

shortfalls, the cosmopsychist still faces a further problem that Micropsychism entirely avoids, namely: a coherently explicated case for championing the anti-reductivism entailed by holism over the reductivism entailed by atomism. Lest we forget that a scientific reductivism, in which upon observing complex wholes it is assumed that the system is non-holistic in such a way that a reductive analysis of its constituent parts will offer a complete understanding of the composite, is the very method that has allowed us to ground an understanding of everything from the double helix to the laws governing classical and statistical mechanics. Thus, contrary to the claims of holists such as Caruana (2000)<sup>36</sup>, the scientific method seems to have progressed perfectly adequately by taking a non-holistic worldview as a postulate, especially given that if holism were true, we would *not* expect a reductive account of a composite's parts to produce a complete understanding of the composite. So, given that there is at least some evidence that a reductivist account of complex wholes is *not* consistently providing us with an incomplete, or faulty, understanding of composites, it seems reasonable to infer that there are, at present, empirical reasons to endorse atomism over holism. Indeed, the lack of empirical evidence for the latter remains one of the strongest arguments in favour of the former, as Chibeni (2004) articulates, 'we have not, a fortiori, any solid scientific or philosophical basis for inferring that the world itself and at large is a single un-analysable whole, or for accepting any other sweeping, unspecific metaphysical statement of ontological holism' (p. 244).

To conclude this section, if we take our awareness of our phenomenal experiences to be evidence of phenomenal experience, and we hold 1) all existent phenomena are spatially and

---

<sup>36</sup> Caruana (2000) argues that if we 'focus exclusively on small units without reference to the wholes they may be part of' we risk a number of 'deformities in our understanding' (p. 5). However, as articulated, it seems that, quite contrarily, an anti-holistic scientific method has *not* consistently generated the sort of 'deformities in our understanding' that would warrant a move to endorse holism *at this time*.

physically extended, 2) phenomenal experience is ontologically irreducible to the spatially extended functions or structures of physical entities, and 3) the perceptual faculties of human subjects cannot be employed to indubitably demarcate the intrinsic nature of spatially extended entities, then the simplest possible explanation for the phenomenon of phenomenal experience seems to be that we simply do not understand the nature of the physical in its totality, and therefore there is enough room in the gap in our understanding to posit phenomenal experience as its intrinsic nature. If we take this articulation of Panpsychist Russellian Monism as our foundation, and we absolve to only commit ourselves to the most explanatorily elegant, or powerful, of the number of variances on Panpsychist Russellian Monism on offer<sup>37</sup>, I suggest that we ought to endorse the least explanatorily taxing version that best fits the evidence and is best able to account for all relevant phenomena. In this case, as articulated within this section, it seems that, if the combination problems can be adequately resolved, the metaphysic with the potential to be the most explanatorily powerful form of Russellian Monism is Constitutive Micropsychism, for this solves the hard problem in a manner that does not rely upon strong emergence to solve the generation problem, does not require that we overlook the lack of empirical evidence for the emergence of novel properties, and does not demand that we face a decombination problem whilst also embracing a paradigm shift away from the atomism that has dominated much of the contemporary scientific zeitgeist. In this respect, if a coherent micropsychist theory of consciousness could be articulated, it seems not unreasonable to hold that micropsychism would be positioned as an elegant solution to the problem of consciousness that can stand as an explanatorily robust option within the field of the metaphysics of the mind.

---

<sup>37</sup> See chapter 3 for an argument in favour of employing explanatory power to delineate between equally naturalised ontological commitments.



This metaphysic remains not without issue, however. Most notably, micropsychism must account for how experiential microphysical entities may be employed to ground a cohesive theory of consciousness that can explain precisely how simple phenomenally unified subjects may constitute macro-level complex conscious subjects (whilst also explaining precisely how micro-level experiential properties do not necessarily manifest ubiquitously at the macro-level). Whilst this thesis endeavours to address all such problems, in order to offer satisfying, explanatorily coherent solutions, it seems befitting to first devote some time to a full articulation of the brevity of these problems that an adequate micropsychist theory of consciousness must address.

### **1.3 The fundamental problems facing Constitutive Micropsychism**

As I write this piece the shrill morning chirps of a chaffinch are accompanied by a wintery breeze that seems to sneak through the jar in the window. These sounds and sensations are unified within a singular subjective state of ‘what it is like’ to exist as this subject of experience at this point in space-time, and it is the fact that there *is* ‘something it is like’ for me to hear the song of a chaffinch, or feel a cool breeze, that affirms my existence most indubitably. Further, these experiences are not isolated and seem to occur simultaneously within a unified experiential locus, in which my experience of the chaffinch’s morning ritual exists alongside my experience of a cold breeze of air. Further still, my uncertainty pertaining to this experience is relatively low in such a way that the ‘informational entropy’<sup>38</sup> (i.e., the

---

<sup>38</sup> For now, it suffices to simply describe informational entropy as a measure of the state of a subject’s uncertainty pertaining to an experience or an event, such that if the informational entropy associated with an experience is high, the amount of available information that may be extracted from this experience is equally high, and subsequently the amount of information the subject holds about the experience is low. Conversely, if the information entropy of an experience is low, the amount of information available to be extracted from said experience is also low because we already hold a large portion of the available information contained within the experience. For example, upon rolling some dice, the informational entropy associated with the event is high

possible informational states) associated with these experiences is relatively low; I can consciously demarcate that the breeze I feel is devoid of the property of heat and I can delineate the sound I hear as resonating from a chaffinch. In this respect, the chirps, the breeze, and the knowledge of their respective properties all occur simultaneously as a conscious symphony orchestrated and unified within a singular phenomenological field of experience that is bound to a singular subject.

It is the beautiful complexity of this conscious symphony that has allowed the human species to attain the higher-level thought necessary to construe the ‘hard’ problem of consciousness, and has allowed certain members of the species to imagine the possibility of this problem finding a solution in bottom-level microphysical entities, such as quarks, holding the ontologically fundamental ‘what it is like’ experiences that eventually constitute my phenomenal experience of a chaffinch’s song on a cold autumnal day. As hinted, however, whilst such a solution would imbue reality with fundamental phenomenological properties and thereby solve the hard problem, in so doing we unavoidably adopt a myriad of deeper explanatory issues which arise as a result of the disparity between the prospect of ubiquitous bottom-level experiential entities and the perceptions we hold of both ourselves and the universe.

In this respect, these explanatory issues can be distinguished into two opposing kinds: 1) the issue of explaining precisely how a multitude of rudimentary, subject-specific microphysical experiences constitute and combine into macro-level, biologically conscious subjects of experience replete with a conscious awareness of their constituent parts, and 2) the issue of

---

(there is a  $1/6$  probability of any particular outcome), but if we were to knowingly toss a weighted dice the probability of a particular outcome increases, and therefore the informational entropy associated with the event decreases. This concept shall be of significant importance throughout this thesis and shall be greatly expanded upon later in the piece (see chapter 6).

explaining how we may reconcile our intuitions about the non-experiential nature of certain physical entities, such as chairs, with the panpsychist commitment to ubiquitous experiential properties. Due to the breadth of these issues, this section shall be split into two sub-sections, the first of which shall address the problem of how subjects constitute other conscious subjects and the second shall address the problem of how ubiquitous micro-experiential subjects may be reconciled with our intuitions about the nature of reality.

### 1.3.1 The Combination Problem(s): Making sense of subjects constituting other conscious subjects

The problem of explaining precisely how entropic, rudimentary micro-experiential subjects may constitute macro-level, organic, conscious subjects seems to contain two sub-problems which must be addressed individually. First, to account for how a multitude of subject-specific microphysical experiences constitute the singular subject-specific conscious experience we hold, we need an adequate solution to the problem of how several individually bounded subject-specific experiences constitute and combine into a new, bounded subject of experience. Second, a solution to the problem of how bottom-level phenomenological properties ground the conscious awareness we hold requires an explicit account of grounding that could make sense of the fact that human and animal consciousness seems to introduce something ontologically ‘over and above’ the qualities and powers entailed within microphysical phenomenological properties. It is the culmination of these two problems that I take to be the reason behind the contemporary move to reject micropsychism, and so whilst the problem of how to reconcile our intuitions about the nature of certain physical entities with micropsychism is of importance (and shall be addressed in depth shortly), I now turn to

offer a deeper elucidation of these two sub-problems. I begin with an articulation of the former.

The combination problem is well known in the contemporary panpsychist literature, and, in its earliest iteration, was construed by James (1890) as follows:

*Take a hundred of them [feelings], 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 101<sup>st</sup> feeling would be a totally new fact; the 100 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, not it with them, and one could never deduce the one from the others, nor (in any intelligible sense) say that they evolved it. (p. 160)*

In this respect, the combination problem is entirely analogous to the decombination problem, for whilst the latter faces the problem of explaining how experientially bounded<sup>39</sup> subjects may split into more experientially bounded subjects, the former faces the problem of explaining how bounded subjects of experience may combine to constitute a further bounded subject of experience. In both cases, if we begin with the private, experientially isolated, subject-specific experience of subject S and either split it or pack it tightly together with other subjects ~ S, we do not produce a new subject, and instead simply either have a myriad of fractured smaller iterations of subject S, or a conglomeration of S and all the other subjects ~

---

<sup>39</sup> As I have done throughout, I take bounded to be synonymous with experientially isolated or subject-specific, in the sense of specifically bound to subject 'S' in a way that is *not* shared amongst subjects ~ S.

S. The point to be laboured is that, in both cases, we have nothing beyond the same experiential subject S, because we are yet to advance a coherent theory for how subject-specific experiences can split or combine. Construed as such, it is not at all clear if we will ever explain how a combination of subjects S and  $\sim S$  could produce a new subject 'K', or how the de-combination of S could ever produce a subject  $\sim S$  because in both cases we must confront the inevitable problem of reconciling the difficulty of subjects constituting other subjects with our deeply held intuition that we each exist as entirely experientially bounded, unique subjects of experience<sup>40</sup>.

Whilst the standard combination problem is certainly forceful and difficult, the problem of explaining precisely how the experiences we associate with biological consciousness can be grounded in the qualities and dispositions of microexperience is arguably just as troublesome as the standard problem of subject-summing. Especially as a solution to this problem must necessarily entail a solution to a myriad of smaller combination problems (i.e., the quality combination problem/the palette problem and the structural combination problem), a solution to the unity problem and the awareness problem (Chalmers 2016, p. 179-184), and must achieve all of this without falling foul to the issues highlighted in Goff's (2018) truth-making account of grounding (an analysis of which is forthcoming). These problems arise because we know that our biological conscious experience involves specific phenomenal qualities such as the coldness of an autumnal wind or the shrill siren of a chaffinch, and we know that these qualities are unified within a singular experiential locus. Further, we know that our experience is structured in a specific way that gives rise to richly complex auditory, visual and spatial fields, each of which are unified and presented to a singular subject. Finally, we

---

<sup>40</sup> As shall be made clear later in the thesis (see chapter 5), for the reasons outlined in chapter 2 I take the subject-summing problem (i.e., how do subjects combine) to be one and the same as the combination problem (i.e., how do experiences combine).

are aware of all of this – indeed, it is precisely our awareness that has allowed us to articulate the hard problem, the combination problem and the myriad of further problems that stem from these, and as a result any coherent panpsychist theory of consciousness simply must explain precisely how microexperiences constitute conscious subjects replete with a consciously unified awareness of both what they are and what their constituents are likely to be. It is our awareness of all of this, coupled with our simultaneous awareness that it is likely that microexperience holds distinct microqualities or microstructures that are not replicated in our biological macro-experience, which has caused Chalmers (2016) to articulate, what he terms, the quality combination problem and the structural combination problem, on the grounds that it seems very likely that the qualities and structures of a quarks experience differ markedly from the qualities and structures of biological experience<sup>41</sup>. As a result, as well as solving the standard combination problem, any adequate panpsychist account of conscious experience must also explain 1) how a relatively simple palette of bottom-level microqualities and dispositions combine to form complex and varied macroqualities and dispositions of the type we associate with biological experience, 2) how such qualities, dispositions and structures become unified within the experiences of a singular conscious subject, and 3) how we, as subjects of experience, achieved the state of awareness necessary to both demarcate ourselves as experiential subjects and articulate the metaphysical problems involved in delineating both what we are and what we are constituted from. Or, put simply, a ‘complete’ panpsychist theory of consciousness must provide an account of how structurally rudimentary, simple micro-experiences constitute subjects with both complexly structured, rich conscious experiences and the conscious awareness necessary to query their ontology.

---

<sup>41</sup> This seems intuitive, for it is highly likely that quarks are unequipped to experience the quality of redness or coldness at all (or at least unequipped to experience them precisely as we do), nor is their experience likely to be structured into orderly visual, spatial, sensory, auditory parameters. So, the problem Chalmers (2016) articulates is one of explaining precisely how the fine-tuned nature of our conscious experiences is grounded in the conceivably highly entropic, simplistic and unstructured experiences of bottom-level microphysical entities.

Whilst my aim in this thesis is to show that an account of this kind is entirely plausible and possible, Goff (2018) argues that because micropsychism is not a form of non-constitutive panpsychism, the micropsychist cannot coherently defend the claim that the biological conscious experience we hold is anything ontologically ‘over and above’ (Heil 2012, p. 63) the qualities and dispositions contained in microexperience, and from here endeavours to construct an argument to show that micropsychism cannot coherently explain how biological consciousness is grounded.

To articulate the flaws implicit in such an argument, I first address Goff’s motivations.

Firstly, I agree with Goff that it is very difficult, if not impossible, to deny that the micropsychist is committed to the contention that biological consciousness is *not* anything ontologically over and above the interplay between physical laws and the interactions that occur between the experiential properties of bottom-level microphysical entities, for to say that biological consciousness is grounded in the more fundamental ontology of microexperiential properties is to say that biological consciousness  $p$  is itself ontologically reducible to microexperiential properties  $q$  in a way that  $p$  introduces into the world  $w$  no ontologically novel causal powers or dispositions that are strongly emergent in a manner that would render them ‘over and above’ the ontologically fundamental properties of  $w$ , and I take such a claim to be implicit to the cogency of constitutive micropsychism. However, Goff (2018) argues that if this is true, the micropsychist is committed to a truthmaking account of grounding, in which we maintain that if a proposition  $p$  is true at a world  $w$ , then  $p$ 's truth at  $w$  is grounded in the fundamental features of  $w$  (Shcaffer 2008, p.10), and therefore we must accept that truths about chairs are only made true by virtue of the ontology of fundamental particles being arranged in a chair-like fashion – much the same as propositions

pertaining to biological consciousness are truth-conditional upon the ontology of fundamental particles being arranged in a certain brain-like way. With this established, Goff (2018) employs the following argument against micropsychism:

1. The metaphysical truth conditions of propositions concerning organic consciousness are a priori accessible.
2. If micropsychism is true, then the metaphysical truth conditions of propositions about organic consciousness concern micro-level conscious entities.
3. It is not plausible that there are a priori accessible metaphysical truth conditions of propositions about organic consciousness which concern micro-level conscious entities.
4. Therefore, micropsychism is false. (p. 11)

The inference here is that if micropsychism holds that the truth conditions for facts about macro-level biological consciousness are grounded in more fundamental microexperiential facts, and if we cannot a priori demarcate micro-experiential properties '*p*' as a necessary truth condition for propositions concerning macro-level conscious entities '*q*', then it is implausible that *p* is a truth condition for propositions concerning *q*, and therefore micropsychism must be false. By this Goff is arguing that if propositions concerning 'what it is like' for an organic subject are only made true by virtue of the macro-level experiences associated with the organic subject in such a way that a proposition concerning macro-level experience cannot have as truth conditions anything more fundamental that is decidedly *not* the macro-level experience of the organic subject, then propositions concerning the 'what it is like' experiences of microphysical entities do *not* act as necessary truth conditions for propositions concerning organic subjects, and therefore micropsychism is false because a truth-making account of grounding cannot account for how propositions concerning macro-



level experiences are grounded in micro-level truths. To compound this argument, Goff relies upon the intuitive inference that when we are concerned with the phenomenal experiences of Bill, we are exclusively concerned with ‘what it is like’ to be Bill and are therefore entirely unconcerned with ‘what it is like’ for the microphysical entities that constitute Bill’s brain. Construed as such, as Goff rightly argues, micropsychism does seem to have a problem with propositions concerning organic subjects such as Bill, because, if micropsychism is true, we are seemingly simultaneously committed to describing Bill’s consciousness as a fundamental ontological truth *and* committed to describing all propositions concerning Bill’s ‘what it is like’ experiences as ultimately only serving to communicate more ontologically fundamental ‘facts about the consciousness of the microphysical entities in Bill’s brain’ (Goff 2018, p.10). Therefore, as fundamental micro-level subjects cannot coherently act as truth makers for propositions concerning fundamental macro-level subjects, if this problem cannot be accounted for, micropsychism simply must be false.

The core of Goff’s sentiments here can be traced to earlier work (see Goff 2015), in which Goff elucidates Sider’s (2009-2012) thesis on constitution, and sympathetically embraces the Siderian notion that:

*‘Truth X constitutes truth Y iff (i) X is a fundamental truth and Y is a non-fundamental truth, (ii) the fundamental reality specified by X satisfies the metaphysical truth condition of Y.’*  
(Goff 2015, p. 373)

Construed as such, we can see that as panpsychism commits us to 1) fundamental micro-level subjects of experience “X”, and 2) *seemingly* fundamental macro-level subjects of experience “Y”, there is no coherent means with which to explain precisely how X may constitute Y. In

this sense, both Sider (2009/12) and Goff (2015) seem in agreement that constitutive forms of Russellian panpsychism commit us to terms such as ‘subjects of experience’ that are ‘jointly carved’ (Goff 2015, p.374) to exist as ontologically fundamental phenomenal truths that simultaneously specify macro *and* micro subjects. As a result, we are left with two fundamental entities denoted by the term ‘subject of experience’: irreducible macro-level subjects and irreducible micro-level subjects. This, as Goff (2015) points out, renders all forms of constitutive Russellian Monism (and therefore Constitutive Russellian Panpsychism) false by committing us to the contradiction that arises upon endorsing 1) the Russellian Panpsychist contention that all macro-truths are grounded, or constituted, in fundamental micro-truths alongside 2) the notion that macro-phenomenal truths are ontologically fundamental.

Whilst this argument appears *prima facie* forceful, I argue that the micropsychist may resist Goff on two fronts. First, it is not at all clear that the truthmaking account of grounding articulated by Goff is the only option available to the micropsychist, and secondly, as premise 2 of Goff’s argument is incomplete, it seems entirely possible to provide an account of constitution that entirely avoids Goff’s (2015/18) and Sider’s (2009/12) concerns. It is sufficient to address the first of these options superficially by simply pointing to the perfectly adequate accounts of grounding that employ the non-truth functional connectivity endorsed by the likes of Hornsby (2005) and Fine (2001/12)<sup>42</sup>, which offer an alternative to accounts, such as Goff’s, that maintain that the truth of propositions such as ‘Bill is conscious’ is

---

<sup>42</sup> Whilst the truthmaking account advanced by Goff is relational in such a way that the proposition that Bill is conscious is true by virtue of the relation to Bill’s consciousness, Fine (2001) argues that the proposition Bill is conscious ‘Y’ is true because of a specific reason ‘X’. The former involves a relation between a truth (Bill’s consciousness) and a truth maker (Bill’s consciousness); the latter involves non-truth-functional sentential connectives.

exclusively dependent upon the relation between the facts A (Bill) and B (consciousness), and instead maintain that the proposition that Bill is conscious is grounded *because* Bill holds some of the qualities we associate with consciousness, or perhaps *because* Bill is constituted of bottom-level micro-experiential subjects that combine to form the organic subject 'Bill'. The crucial difference between these two accounts is that the former grounds the truth of propositions in the truth of a proposition's constituent facts (i.e., the proposition 'Bill is conscious' is true by virtue of the fact that Bill *is* conscious), whilst the latter simply grounds propositions by offering reasons for why they should or should not be believed, and therefore offers an account that remains neutral on the issue of what precisely grounds, or does not ground, the proposition that Bill is conscious. In this respect, if we were to endorse this latter account of grounding, we would be able to avoid Goff's argument entirely because we would be able to ground the proposition 'Bill is conscious' by employing the non-truth functional connective 'because' to ground the proposition as follows: 'Bill is conscious because micro-level conscious subjects combine to form Bill'.

With this established, I move on to address my second, arguably more forceful, means of resisting Goff's argument against micropsychism. This point of contention allows us to maintain the advantages of the truthmaking account of grounding endorsed by Goff by simply highlighting the incomplete nature of Goff's (2018) second premise:

2) If micropsychism is true, then the metaphysical truth conditions of propositions about organic consciousness concern micro-level conscious entities.

Here, Goff is maintaining that the truth conditions of propositions concerning organic consciousness are grounded, *in entirety*, in micro-level conscious entities, and seems to imply

that this is something that the micropsychist cannot refute. However, quite contrarily, micropsychism is not committed to the notion that macro-level organic consciousness is *wholly* grounded in micro-level conscious entities. As Chalmers (2016) notes, whilst many philosophers of mind (including himself) have framed the problem as one of: ‘how do microexperiences come together to yield X’, or ‘how do microexperiences ground macroexperience’, it is in, fact, incorrect to frame the issue as such. Instead, it is entirely possible for macroexperience to ‘be partly grounded in causal or structural relations among the microexperiences, or other microphysical properties, or even in other quiddities if there are non-phenomenal quiddities as well’ (Chalmers 2016, p. 184), and as a result, we should say that ‘constitutive panpsychism requires macroexperiences to be wholly grounded in microexperiences *and* microphysics, where microphysics is understood broadly to include all of the above [emphasis added]’ (ibid, p. 184). As such, Chalmers (2016) argues that we should frame the fundamental issue for micropsychism as one of accounting for how ‘microexperiences and microphysics come together to yield X’ (p. 184), and I would agree. With this established, and in lieu of a specific argument to address precisely why macroexperience must *exclusively* be grounded in microexperience, it seems reasonable for Goff (2018) to allow for the possibility that his second premise is incomplete, and subsequently should be reformulated as follows<sup>43</sup>:

2) If micropsychism is true, then the metaphysical truth conditions of propositions about organic consciousness concern micro-level conscious entities *and* microphysics.

---

<sup>43</sup> I note that in places (see Goff 2016), Goff does seem open to the notion that we might explain the existence of macro subjects by appeal to the relations between micro parts. He does note, however, that this seems to lead us to a kind of mysterianism in which the nature of this relation is rendered vague. I shall address this particular concern later in the piece (see chapter 5).

If this is the case, as Chalmers (2016) argues, the micropsychist has far more explanatory resources at their disposal and, as I shall endeavour to show in this thesis, is equipped to address Goff's (2018) call for an alternative account of grounding the likes of which has been sorely missing from the contemporary panpsychist literature:

*'Perhaps there is some way of construing micropsychism other than the truthmaking account I have given above. However, this would require formulating some other account of grounding, an account that incorporates some other way of making sense of organic conscious minds being 'nothing over and above' facts about micro-level minds. I have found nothing like this in the literature so far.'* (Goff 2018, p. 11).

Crucially, to provide such an account of grounding, we must explicate a form of constitution that begins with instantiations of fundamental microphenomenality and microphysics "X" and ends with an instantiation of macrophenomenality "Y" that is nothing ontologically over and above the properties and causal powers of "X". Construed as such, we reach the final means of resisting Goff's rejection of micropsychism: providing a means of constitution, or grounding, that can explain how X constitutes Y in a manner that explains how Y *is* a real existent and *is* nothing ontologically over and above X<sup>44</sup>. To achieve this, we must uphold the central tenet of constitutive panpsychism by maintaining that macro-subjects do not introduce any new causal powers or properties that are irreducible to the micro-level constituents and microphysics whilst also explaining how microphysics and microexperiences (X) may constitute macrophenomenal facts (Y), if 1) Y obtains in virtue of X and 2) Y is nothing ontologically over and above X. If such an account can be achieved, it would entirely avoid

---

<sup>44</sup> In this respect I concede that Y might satisfy what Armstrong (1997) calls an 'ontological free lunch' (p. 12), but note that in the coming chapters I will endeavour to offer an account of how, in this case, we might achieve this 'free lunch' (see chapters 4/5) without inviting any unsavory ontological commitments.

Goff's (2015/18) and Sider's (2009/12) concerns by virtue of positing that Bill's macro-level consciousness does not stand as a fundamental subject *in addition to* the fundamental subjects that constitute it because it *is* nothing ontologically over and above a complex manifestation of the microexperiences and the 'causal or structural relations among the microexperiences' (Chalmers 2016, p. 184). However, a thesis of this kind must overcome an array of explanatory hurdles, as Goff (2015) elucidates:

*What is it, though, for a certain fact to be nothing over and above another fact? Reflection on crowds, parties or organisations makes the notion intuitive. But philosophical reflection can render it somewhat mysterious. A crowd is neither identical to its members, nor wholly distinct from them. What is this strange middle way between identity and distinctness? How can fact X involve different objects and properties to fact Y, and yet, from the perspective of serious metaphysics, add nothing beyond the objects of properties already involved in Y? Philosophers trading in 'nothing over and above' talk owe us an account of how they get their lunch for free. (p. 382)*

Parts of this thesis could be construed as a direct response to Goff's sentiments in this passage, and in the subsequent chapters (see chapters 5/6) I will endeavour to provide precisely the sort of account that Goff is calling for. First, however, after having articulated what I termed within my introduction to this section as the first problem facing a micropsychist theory of consciousness, I now move on to articulate the second problem facing an adequate micropsychist theory of consciousness: the issue of explaining how ubiquitous phenomenal properties may be reconciled with our intuitions about the seeming non-phenomenal nature of certain physical entities.

### 1.3.2 The ‘Inverted’ Combination Problem: reconciling Micropsychism with our intuitions

Whilst it seems to be something of a patent truth that our perceptual faculties are limited in such a way that leaves room for the possibility of Russellian micropsychism, it remains equally true that micropsychism must be reconciled with our common-sense intuition that inanimate physical objects, such as rocks or chairs, are *not* experiential. This reconciliation is of particular concern because it appears that, in this case, our intuition seems to typify a compelling reason to reject micropsychism outright; we have no empirical justification for the inference that rocks, for example, are experiential, such physical entities do not seem to react to the world in the same way experiential agents do, and therefore adopting any form of panpsychism seems to entail the adoption of an unintuitive, and empirically unverified, inference.

The problem with this argument, however, is that micropsychism does not necessarily commit us to experiential rocks or chairs. Quite contrarily, micropsychism only commits us to the experiential nature of bottom-level microphysical entities, and so we may conceivably establish a means with which to explain how certain conglomerations of microphysical experiential entities may give rise to experiential complex wholes of the type we witness in conscious biological agents whilst other conglomerations may give rise to certain non-experiential wholes, such as rocks or chairs. Basile (2010) articulates this argument as follows:

[A] All ultimate constituents of reality are sentient

[B] When such ultimates are grouped in certain special ways – say, the experiential ways – they give rise to complexes that are themselves sentient

[C] When they are grouped in different ways – say, the non-experiential ways – they give rise to complexes that are themselves insentient. (p. 98)

As Basile (2010) points out, we commit ourselves to the fallacy of composition if we hold that a complex whole must necessarily be experiential if its parts are, just as we commit ourselves to the fallacy of division if we hold that a non-experiential whole necessarily reveals non-experiential parts (p. 98). So, it is entirely plausible that certain conglomerations of experiential microphysical entities could give rise to both experiential and non-experiential wholes, and therefore *prima facie* it seems that micropsychism does not necessarily condemn us to uphold the unintuitive inference that inanimate objects are experiential.

However, to coherently avoid this unintuitive inference, the micropsychist must articulate an explanation for how phenomenal properties manifest at the micro-level but do *not* manifest in certain macro-level conglomerations of microphysical entities. Simply, the micropsychist needs an adequate account for how phenomenal properties are occasionally ‘lost’ at the macro-level. This is what I call the ‘inverted’ combination problem<sup>45</sup>, i.e., the problem of explaining precisely how certain conglomerations of experiential microphysical entities do *not* combine to form macro-experiential entities. Basile (2010) describes this as the problem of accounting for which ‘principle’ causes only certain organisations of microexperiences to produce a unified experiential subject and maintains that this principle must coherently distinguish between ‘those ways of organisation that give rise to sentient wholes, and those

---

<sup>45</sup> So called because it seems to present the opposite dilemma to the issue facing the more standard articulation of the combination problem. The combination problem requires an explanation for how microexperiences combine to form macro-subjects; the inverted combination problem requires an explanation for how experiences manifest ubiquitously at the micro level but do *not* manifest ubiquitously at the macro level, or more simply, an explanation for how microexperiential properties combine to form anything other than macroexperiential properties.



that do not' (p. 98-99). Whilst I am broadly in agreement with Basile's articulation of the 'inverted combination problem', I note that the core of this issue is accounting for how phenomenal properties manifest ubiquitously at the micro-level but do *not* manifest ubiquitously at the macro-level. As Basile (2010) notes, the panpsychist literature is currently lacking a thoroughly articulated solution to this problem, and, as a result, it seems that most previous panpsychist theories of consciousness struggle to adequately explain away the unintuitive inference that inanimate objects are experiential.

With this established, I am now in a position to offer a simple articulation of the two problems that must be addressed within any adequate panpsychist theory of consciousness: 1) an explanation for precisely how simple, bottom-level microexperiential entities combine to form a singular, bounded, conscious subject with complex, rich phenomenal content/structure and an awareness of itself and its constituent parts, and 2) an explanation for how the phenomenal properties of bottom-level microphysical entities may manifest ubiquitously at the micro-level but *not* necessarily manifest ubiquitously at the macro-level. Articulated as such, this thesis is concerned with answering the fundamental question: *how do microexperiences and microphysics ground both richly consciously aware macro-experiential subjects and non-experiential macro-level entities*. In what follows, I offer a Russellian Micropsychist theory of consciousness with the potential to provide a solution to this problem in a manner that does not violate our current scientific worldview or resign us to any form of supernaturalism. This solution takes us from the disposition(s) of bottom-level microexperiential entities to the possibility of construing a model for a 'Phenomenal Maxwellian Demon' capable of sustaining itself away from thermodynamic equilibrium long enough to produce a far from equilibrium system that competes with other such systems until thermodynamic natural selection guides us to an account of human conscious experience that

is perhaps capable of resolving some of the problems afflicting Russellian Micropsychism. To reach this point, however, we must first plug the hole in our understanding of the physical by shedding some light on ‘what it is like’ to be a bottom-level microphysical entity, for it is only upon properly conceiving of the nature of bottom-level phenomenal properties that we might begin to reveal how our consciousness arose from these simple constituents.

## Chapter 2

### On the nature of phenomenal properties

While there are various disputes pertaining to the limits of both our experimental and epistemic capacities to indubitably demarcate the nature of bottom-level entities (See Aad et al 2015, Fox 2009), if we maintain that there *is* a bottom-level, it would be a feat of some relatively subversive conviction to deny that our current best guess at what might be described as *the* fundamental microphysical entity is: the quark<sup>46</sup>. According to the ‘standard model’ currently dominating contemporary physics, such quarks are the indivisible foundations of reality constituting 99.9% of all the perceptible matter in the known universe, and as minute specks of energy with a diameter of a mere million millionth of a millimetre ( $10^{-15}$  metre), they orbit at near the speed of light inside the protons and neutrons they constitute, in a sort of micro-level circadian rhythm that acts as the bedrock of our physical reality. Regardless of whether the physical sciences are right to posit quarks, or leptons, as the fundamental constituents of reality, for the reasons advanced in the previous chapter, it remains the case that 1) atomism probably, in lieu of evidence to the contrary, offers our most explanatorily coherent account of reality and is, again, in lieu of evidence to the contrary, justified in the notion that there *is* a bottom-level microphysical entity that constitutes reality as we know it, and 2) in order for this bottom-level microphysical entity to be reconciled with

---

<sup>46</sup> Quarks are in fact one of two categories of entities that the standard model takes to be indivisible and thereby fundamental. The other is leptons. Further, I note that there are multiple alternatives to the standard model which might dispute the claim that quarks are ‘the’ fundamental microphysical entity - one alternative is string theory, which replaces talk of point-like particles (quarks) with one-dimensional ‘strings’. With this mentioned, I note that I do not devote any time to an analysis of these (or any other) positions here. Instead, I take it that as all scientifically credible theories of fundamental reality posit spatially extended constituents with structural properties of some kind, from the perspective of Russellian Micropsychism it does not matter which instantiation of physicality is demarcated as fundamental, as long as this instantiation is intrinsically phenomenal. In this respect, I hold that whichever instantiation of physicality turns out to be *the* fundamental instantiation just *is* intrinsically phenomenal. With this in mind, for simplicity I frequently refer to quarks as ‘the’ fundamental entity, but in actuality quarks here might often stand for ‘whichever entity turns out to be fundamental’.

our existence as conscious subjects, it is explanatorily advantageous to endorse a form of Russellian Micropsychism by positing that this bottom-level entity instantiates *both* micro-level spatially extended structure *and* micro-level phenomenality.

With this established, the problem is now one of comprehending precisely what micro-phenomenal properties are, for to picture the world in terms of microphysics renders it nearly incomprehensible to the human mind. Imagine, for example, shrinking ourselves down to a millionth of a millionth of a millimetre, such that we may perceive with clarity the quarks that ostensibly constitute 99.9% of the perceptible macro-level matter. It would be immediately clear that it is very likely that the experiences of these quarks are, for the most part, incomparable with the experiences of the macro-level brains they constitute, and yet the experiences associated with macro-level human brains are the only experiences we are epistemically privileged enough to make inferences about. How, therefore, could we ever reasonably infer ‘what it is like’ to be a bottom-level microphysical entity? In answer to this, I posit that we must begin at the macro-level of human conscious experience and from there sharpen our concept of phenomenal properties to such an extent that we reveal a coherent model for such properties that can retain internal consistency, and be applied to bottom-level microphysical entities, without contravening our understanding of the properties as they manifest at the macro-level of human conscious experience. In this sense, while it is clear that upon discussing phenomenal properties, we are broadly discussing the (or ‘a’) property of *subjective conscious experience*, a full account of precisely what subjective conscious experience *is*, or how this property might manifest at the bottom-level, remains elusive without a full account of what precisely the qualities (and/or dispositions) of phenomenal properties broadly construed are.

I, therefore, suggest that to grasp the intrinsic nature of the physical and provide a fitting foundation from which to begin to address the problems afflicting constitutive micropsychism, we must first offer an explicit account of precisely what we are referencing when discussing phenomenal properties. To present a thorough account of this kind, I maintain that we must provide an outline for both 1) the necessary qualities of phenomenal properties, i.e., the qualities that are essential to the formation of a coherently and completely realised concept of phenomenal properties, and 2) the role these phenomenal properties may play in the natural world<sup>47</sup>. With this aim in mind, contrary to Rosenthal (2001/05), I posit that whilst introspection alone cannot be employed to indubitably delineate the definite qualities of phenomenal properties, it remains the best possible starting point in our endeavour to grasp the nature of such properties because it is only by virtue of our subjective conscious experiences (or, more precisely, our instantiating phenomenal properties) that we are aware of our phenomenal states and thereby their qualities at all. In this vein, we must begin with an analysis of ‘what it is like’ to *be* a macro-level conscious subject, and therefore I begin with a reflection on my own conscious experiences.

Having become accustomed to the shrill morning chirps of the chaffinch that had been sharing its song with me of late, on this morning I notice a sudden queer silence permeating the content of my experiential state, and surmise that the chaffinch has seen fit to share its song with pastures new. In this ensuing quiet, I, as the subject of experience, experientially demarcate the differences between the qualitative content that manifests at this moment in space-time and that which occurred before, and this shift in my experience of ‘what it is like’ for me as a conscious subject seemingly enables me to employ phenomenal content to make

---

<sup>47</sup> An account of the role phenomenal properties play will be hinted at in the next section but will be thoroughly expanded upon in chapters 3 and 4.

inferences pertaining to the physical properties of my environment. From this realization alone, we may demarcate at least four basic qualities and/or dispositions that seem to be essential to the phenomenal properties that manifest at the level of human consciousness: 1) when we engage in discussions pertaining to the nature of phenomenal properties, our phenomenal properties seem to be having a marked causal impact upon the physical behaviour we exhibit, 2) the intelligibility of the concept of phenomenal properties seems to rely on an *experience-for* relation and thereby entails a subject to whom phenomenal experiences may be ascribed (a subject of experience), 3) phenomenal properties seem to be, at least partially, transparent *experiences-of* (in the sense that they seem to be experiences of content of some kind), and it seems *prima facie* that this transparency holds the potential to reveal something about the nature of mind-independent physical properties, and 4) being a subject of experience seems to confer the power to individuate between distinct types of phenomenal content. If this account holds, we seem to have presented three qualities that are implicit to phenomenal properties and at least one potential power that may arise as a result of such qualities.

In the name of brevity, I shall devote this chapter to an account of the three qualities that I take to be implicit to phenomenal properties and shall cover an account of their potential powers when I come to find a place for these qualities within the natural world in chapters 3 and 4<sup>48</sup>. The task of this chapter is therefore to delineate whether the qualities I have outlined *are* implicit to phenomenal properties and delineate which may be stripped without rendering the concept of phenomenal properties unintelligible. To achieve this, this chapter shall initially be split into three sub-sections. The first shall address the claim that phenomenal

---

<sup>48</sup> See p. 166 for this account.

properties should be thought of as causally relevant, the second the extent to which phenomenal properties are necessarily ‘experiences-for’, and the third the extent to which we may accurately describe phenomenal properties as contentful and transparent. With this established, I hope to be able to offer a tentative theory of ‘what it is like’ to be a bottom-level microphysical entity that may serve to plug the hole in our understanding of the physical and provide us with a foundation from which to address the problems facing constitutive micropsychism. I begin by addressing the causal efficacy of phenomenal properties.

## **2.1 In favour of Powerful Qualities: On the incompleteness of the physical account of causation and the causal relevance of phenomenal properties**

The zeitgeist of organized philosophical and scientific thought is, for the most part, committed to the contention that the structural relations analysed by physics can provide a fully realized and complete account of causation. The standard account therefore resigns consciousness to a causally inert epiphenomenon that is causally insignificant to the structural properties and behaviours we witness in the natural world. In this sense, there is a tendency to exclude mental, or phenomenal, causation from our scientific worldview on the grounds that 1) ‘no single event can have more than one sufficient cause occurring at any give time’ (Kim 2005, p. 42), and 2) all physical events are uniquely accounted for by other physical events. This argument may be articulated as follows:

1. The mental is not identical with the physical (mental distinctness)<sup>49</sup>.

---

<sup>49</sup> As we shall see, from the perspective of Russellian Micropsychism we might overcome this premise relatively easily by suggesting that phenomenality *is* identical to physicality.

2. Physical events have sufficient physical causes if they are caused at all (physical adequacy).
3. Therefore, no physical events are uniquely caused by non-physical events. (Kim 2005)

The problem with this is that it seems entirely reasonable to define a property ‘X’ as exemplifying causal significance if X’s existence increases the possible states the universe may occupy<sup>50</sup>, and therefore by engaging in discussions pertaining to phenomenal properties, it appears not unreasonable to suggest that these properties are displaying at least some sort of causal significance. Further, we might even suggest that in partaking in discussions pertaining to the unreality of causally efficacious phenomenal properties, the advocate of consciousness epiphenomenalism invariably self-stultifies upon attempting to deny the causal efficacy of the property in discussion. This is because if phenomenal properties truly are epiphenomenal, they should have no causal impact upon the universe at all and thereby should exist beyond the scope of things to which we may have the knowledge necessary to stand in a propositional attitude towards. Therefore, given that 1) we cannot easily deny the causal significance of phenomenal properties without self-stultifying, and 2) a sufficient model for reality *must* account for all known phenomena in such a way that a model predicated upon the causal closure of the physical seemingly cannot, it seems not unreasonable to infer that the (narrowly) physical account of causation is at best incomplete, and as a result, we ought to revise our concept of causation accordingly. To substantiate this sentiment, in what follows I analyse and defend the self-stultifying principle from the contemporary attacks erected by Robinson (1982/2013) and argue in favour of the claim that

---

<sup>50</sup> Rosenberg (2004) employs the same definition of causal significance.



our knowledge of phenomenal properties is itself causal in such a way that invariably reveals the causal efficacy of phenomenal properties. From here, I move to articulate the core of the conceptual issues underpinning the epiphenomena debate, before offering, in the vein of Goff (2017/19) and Chalmers (2013), an account of phenomenal properties that avoids phenomenal epiphenomenalism by placing them as a categorical property of the physical. Finally, I present Robinson's (2014/18) and Howell's (2014) recent moves to call into question the coherence of employing an account of Russellian Panpsychism of the form advanced by Goff (2017/19) and his contemporaries, and subsequently defend Russellian Panpsychism by adopting a Heilian (2015) account of powerful qualities.

### 2.1.1 The Self-Stultifying Objection to Epiphenomenalism

As articulated, the self-stultifying objection perhaps encompasses the most compelling reason to reject phenomenal epiphenomenalism outright, for the epiphenomenalist cannot coherently hold that *all phenomenal properties are causally insignificant* without inviting the contradiction that occurs upon the epiphenomenalist claiming that there *is* a phenomenal property 'X' such that X holds the property of causal insignificance. Framed as such, the epiphenomenalist 'is caught in the self-stultifying position of reporting that we have qualitative events, but at the same time having no justification for the report that we have qualitative events' (Moore 2012, p. 630), by virtue of simultaneously maintaining 1) X exists alongside 2) X is causally insignificant. This is problematic because, in this case, if the latter is true, the justification for the former is lost on the grounds that we cannot reasonably justify the formation of inferences pertaining to truly causally insignificant properties.

Subsequently, it seems that the epiphenomalist must either deny that Xs exist at all, but in so doing must arguably forego their epiphenomenalism by conceding that they are incapable of ascribing the property of causal insignificance to that which purportedly does not exist, or maintain that Xs exist whilst denying their causal significance, but in so doing commit themselves to the self-stultification that occurs upon ascribing causal insignificance to a property that, by virtue of our standing in this propositional attitude, is clearly expanding the possible states the universe might occupy. The strength of this objection is largely recognised within contemporary philosophy of mind, with various philosophers calling to reject any narrative in which the focus is *whether* phenomenal (or mental) properties are causally significant, and instead focus on narratives that can expound coherent metaphysical frameworks for precisely *how* our conception of causation can be revised in such a way that leads to the reconciliation of mental, and as a consequence phenomenal, causation<sup>51</sup> with our scientific beliefs (for proponents of this view see Kim 1998, Robb & Heil 2014, Heil 2012). However, there are also those (see Robinson 1982/2006/2013) who doggedly rebut the strength of this self-stultifying objection and thereby steadfastly deny the call to revise our concept of causation on the grounds that phenomenal properties stand related to the physical body ‘as the bell of a clock to the works, [such that] consciousness answers to the sound which the bell gives out when it is struck’, but holds no bearing upon the behaviour of the workings (Huxley 1874, p. 571).

---

<sup>51</sup> I note that here although I am conflating mental and phenomenal causation, the scope of my enquiry in this section extends only so far as to offer sufficient reasons for why phenomenal properties are themselves causally relevant. As phenomenal consciousness is typically taken to be a robust candidate for at least one ‘mark of the mental’ (see Tartaglia 2008, Pernu 2017 amongst various others), and as the same arguments that have been erected to substantiate mental epiphenomenalism can also be erected to substantiate phenomenal epiphenomenalism, I take it that in this case the conflation is not ill-founded, especially as the solution I advance to the problem of phenomenal causation might reasonably also apply to the broader problem of mental causation.

Robinson (1982-2018) maintains that our knowledge of phenomenal properties is itself no justification for the notion that phenomenal properties are necessarily causally significant, on the grounds that neural states cause phenomenal events (Robinson 2018, p. 10-12), and therefore it is exclusively the neural states themselves that are causally significant. Robinson is keen to demarcate this model for what he describes as a ‘Qualitative Event Realism’ (QER), which posits that qualitative events (subjective experiences) are *caused* by physical structures, from the more standard Identity Theory, which posits that the subjective conscious experience of the blackness of these letters, for example, is a physical event *constituted* exclusively by the relations between the physical properties of neural mechanisms. The argument follows that whilst a property *constituted* by physical mechanisms ought to exclusively manifest physical properties, the same need not necessarily be said for properties *caused* by physical mechanisms, for ‘properties of effects need not be found in the causes of their effects, and conversely’ (Robinson 2018, p. 17). In this sense, whilst phenomenal events may be caused by the physical properties of neural states, as they are themselves simply events that arise as the effects of their physical causes, they may coherently, according to Robinson (2018), instantiate novel properties that are not ontologically reducible to the lower-level domain from which they emerged. Subsequently, whilst we might believe that we hold causally significant phenomenal properties, and we might believe that said properties are entirely irreducible to their physical counterparts, at core these properties are simply the properties of inert events and *not* properties of actual, existent entities (i.e., properties of quarks, or brain states). Therefore, in a not wholly dissimilar vein to Dennett (1991/3), Robinson seems content to position our belief that brains hold causally significant phenomenal properties as nothing beyond the illusion caused by the brain causing phenomenal events, which are epiphenomenal by-products of exclusively *narrowly* physical neural relations. Construed as such, Robinson ostensibly avoids the problem of self-

stultification by maintaining that we can explain and reference phenomenal events without immediately surmising that causal efficacy extends beyond the physical.

The problem with this view is twofold. Firstly, it is not immediately clear that, when appropriately contextualised, ‘properties of effects need not be found in the causes of their effects’ (Robinson 2018, p. 17), for whilst this general principle might be true of physical causes and effects, it is less clear that this makes sense when discussing the non-physical effects of physical causes. This is because whilst, in the case of physical effects/causes, there might very well be new properties, it remains the case that these properties are ontologically reducible to their physical causes – that is, whilst liquidity is a novel property caused by loosely bonded H<sub>2</sub>O molecules sliding past one another, this novel property is *not* ontologically novel in the strict sense because it is still itself physical and is, therefore, at least in principle, deducible from and reducible to an underlying physical cause. Conversely, in the case of non-physical effects of physical causes, it appears that we are necessarily discussing the ontologically novel properties of non-physical effects that are strictly not deducible from or reducible to the physical properties of their causes, and therefore it seems in order to explain how a non-physical effect can be brought about by a physical cause Robinson (2018) needs to rely on the possibility of strong (ontological) emergence. This, as I’ve articulated in the previous chapter, is considered a largely undesirable commitment, for we are currently missing a detailed account of precisely how we can overcome the inconceivability issues associated with the contention that irreducible, novel, non-physical properties can causally emerge from a lower-level narrowly physical phenomenon that does not in any way ontologically necessitate or explain the existence of such properties. Further still, in committing himself to this particular claim, Robinson seems to undermine the central motivation for embracing epiphenomenalism, for if we simultaneously uphold the intuitive

notion that causal relations require a common nexus (such that they share a commonality within which cause and effect may interact<sup>52</sup>), and we embrace Robinson's thesis that certain non-physical events are uniquely caused by physical events, then Robinson seems to face precisely the interactionist issues espoused against substance dualism<sup>53</sup>. This is because delineating the precise nature of a common nexus that connects the distinct ontological identities of the physical and the non-physical remains extraordinarily difficult (especially if we maintain a *narrow* physicalism that embraces causal closure), and as a result, it is not at all clear how narrow structural properties may ever stand in the right causal relation to act as the cause of non-physical events. As Robinson gives us no reason to believe that he is an anti-realist about causation and persistently upholds the possibility of physical neural relations directly bringing phenomenal events into being, then it seems not unreasonable to posit that Robinson is committed to the contention that a physical/non-physical causal nexus *does* exist, and thereby seems to invite the question: if the physical employs a common nexus to cause the non-physical, why can the non-physical not employ this same nexus to cause the physical?

Arguably, contrary to his aims, Robinson (2018) offers an answer to this question by virtue of failing to convincingly avoid the self-stultification principle, for it seems that even if we accept the thesis that a purely physical neural relation (N1) causes phenomenal event 'M' at time 1 (T1), before, at T2, N1 proceeds to cause neural states N2, N3 et cetera, which culminate in the mistaken belief that 'M' exists as an ontologically distinct, causally

---

<sup>52</sup> See Bunge (1982) and Gibran (2014) for articulations of why a causal nexus is perceived by many to be a necessity when discussing causal relations. I note that Robinson could conceivably reject the causal nexus thesis, but in so doing would face a difficulty in providing an account of causation that does not require a commonality within which cause and effect may interact.

<sup>53</sup> Specifically, the causal issue of accounting for precisely how a non-physical entity may cause physical effects.

efficacious phenomenal property of an *entity* instead of the belief that M is simply the epiphenomenal property of an *event* caused by N1 at T1, it seems ‘M’ necessarily still exemplifies its causal significance as soon as we make any inference at all pertaining to the nature of M, because the inference itself has expanded the possible states the universe may occupy in a direction that conceivably would not have occurred in the absence of M, or if M were truly epiphenomenal. In this respect, Robinson seemingly cannot escape the contention that phenomenal properties are causally significant, for even if we overlook the explanatory shortfalls of describing them as non-physical effects of physical causes, upon entering into any description of these non-physical effects at all, the epiphenomenalist invariably self-stultifies by virtue of forming a belief about that which, if epiphenomenalism were true, should be beyond the scope of things which we may reasonably make inferences about<sup>54</sup>.

Whilst this may be the case and may provide a reasonable response to those who accept that *something* is causing our belief in phenomenal properties (but reject that the nature of this cause is phenomenal), the argument I have provided is still conceivably subject to objections centred on the contention that I am conflating causal significance with causal efficacy. On such an account (see Nagasawa 2010), it might be argued that whilst phenomenal properties are perhaps causally significant, our knowledge of them may be grounded in a non-causal relation (perhaps some form of direct self-acquaintance) and thereby the self-stultification argument would not convincingly justify an appeal to their causal efficacy. This would be the

---

<sup>54</sup> At this stage, the only option available to Robinson (2018) is to persist in the contention that phenomenal events do not expand the states of the universe at all by falling back onto a far more extreme form of illusionism, in which, like Dennett (1993), he would be forced to posit that phenomenal properties are entirely illusory and wholly non-existent (and in so doing invite the contestation: ‘illusory to whom?’). Whilst this would perhaps somehow allow Robinson to maintain that the perceived causal efficacy of phenomenal properties is an illusion by maintaining that phenomenal properties are themselves an illusion, this seems fundamentally at odds with Robinson’s aim of simply highlighting the illusory nature of our belief that phenomenal properties belong to entities instead of events, and arguably, even if this extreme illusionism were endorsed, Robinson would still adopt all of the objections mounted against Dennett 1993 (see Chalmers 1996, Carruthers 2005 for a survey of such objections).

case because whilst I have upheld that causal significance refers to any instance in which an entity's existence increases the possible states the universe may occupy (to the extent that an entity 'X' is necessarily causally significant if its existence increases the states the universe may occupy), the existence of X's causal significance is not necessarily in itself proof of X's causal efficacy. It might even be tempting as a result of this inference to argue that X perhaps exists epiphenomenally whilst also increasing the states the universe might occupy (although, as I shall argue, given the criteria I am using this seems like somewhat of a contradiction), and this argument would fit well with those who might be sympathetic to the claim that our belief in, or knowledge of, a phenomenal property 'X' is grounded non-causally, such that X is not causally efficacious but *is* causally significant.

In response to this, I suggest that if our knowledge of 'X' is not caused by 'X', then the advocate of the above position faces a substantial difficulty in substantiating where our knowledge of 'X' came from. As De Brigard (2014) highlights, whilst Robinson (1982/2006/2013) aims to weaken the self-stultification argument by emancipating us from the belief that our knowledge of our sensations must necessarily be caused by the sensations themselves, the consequence of this is that epiphenomenalism must embrace the explanatory burden of accounting for precisely where the knowledge of the sensations we reference originated.

To render the nature of this problem transparent, I note that it appears not unreasonable to hold that to *know* about the nature of a phenomenal event 'X', one must know what 'X' is about, or 'like', in the phenomenal sense, such that our knowledge of X is not merely some retrospective intentional content held in the mind, but a direct association with 'what it is like' to be in a given state 'X'. Now, even if a given phenomenal event *is*, for the sake of

argument, itself taken to be epiphenomenal in the sense that it is a mere non-causal by-product of our direct self-acquaintance with a particular brain-state, then it remains the case that any future brain state that references the event *in full* cannot help but employ the phenomenal knowledge of what the event was initially about, or phenomenally 'like'. Yet, if this knowledge is not causally connected to the initial event in any way whatsoever, we might press the recipient of the event for an account of precisely where this phenomenal knowledge originated if it was not caused by the event itself. And more, if this knowledge is truly not causally connected to this initial event, and thereby not phenomenal knowledge of 'what it is like' in the strict sense, how can one justify the claim that this is knowledge in the relevant sense<sup>55</sup>?

Gertler (2019) argues that this particular concern can be 'assuaged by the claim that we can grasp phenomenal properties through acquaintance', because 'in knowledge by acquaintance, the phenomenal properties serve as the mode of presentation for their referents' (p. 74). By this Gertler means that phenomenal properties do not *cause* phenomenal knowledge of any kind, but instead, our phenomenal knowledge is contained in the 'modes' in which phenomenality is presented to us. In the sense that upon consuming an unsavoury food source we might achieve phenomenal knowledge that this referent is foul, but the knowledge would not be caused by phenomenality, it would merely be phenomenal knowledge achieved via a direct (non-causal) acquaintance with the mode of presentation associated with this instantiation of phenomenality. The problem with this is that all this serves to do is push the

---

<sup>55</sup> In the sense that if we have a myriad of brain-states S1, S2, S3 et cetera that are causally connected in the direction of a report that advances knowledge of a given phenomenal event 'X', if this phenomenal event is not admitted into the causally efficacious brain states that produce this report, how can the knowledge of this event ever been coherently framed as knowledge in the sense relevant to phenomenally knowing 'what it is like' to be in, or undergo, phenomenal event 'X'. Framed as such, I take it that as we do seem to have phenomenal knowledge of the type I discuss (i.e., we *can* retrospectively recall or know 'what it was like' to undergo a given phenomenal event), then this might stand as a sufficient reason to reject the reality of epiphenomenal phenomenal properties all together.



causal ground for our retrospective phenomenal knowledge back to more fundamental states of self-acquaintance, for regardless of whether the phenomenal properties are modes of presentation or not, it seems that as soon as we retrospectively reference a phenomenal event in full (i.e., in a way that might also capture the phenomenal properties associated with the way in which we were initially acquainted with the event), we cannot help but admit the causal efficacy of at least some of the phenomenal properties associated with this event into the causal chain that led to the epistemic state we are in when referencing the event. In this case, because phenomenal knowledge requires (or *is*) a phenomenal acquaintance that entails ‘having the relevant phenomenal experience’ (Gertler 2019, p 64), it does not particularly matter whether we demarcate the ostensive lack of causal efficacy associated with the event, or its phenomenal properties, from the causal efficacy associated with our knowledge of the event because it appears that the knowledge of ‘what it is like’ to have undergone the event is inseparable from the event itself (in the sense that properties of the latter is an intrinsic feature of the former), and therefore upon retrospectively consulting our phenomenal knowledge of having been in a given phenomenal state we also unavoidably reference the associated phenomenal event ‘X’ (and its associated phenomenal properties). Therefore, it is entirely unreasonable to attest that ‘X’ is epiphenomenal because our direct self-acquaintance with ‘X’ is still an intrinsic feature of the phenomenal knowledge that has caused a specific type of future epistemic state (i.e., a state in which we retrospectively know ‘what it is like’ to have undergone this event) that would not have reasonably occurred in the absence of X (or if X were epiphenomenal).

I subsequently hold that my conflation of causal significance and causal efficacy is not obviously misguided, because, in this context, these are one and the same thing. I am arguing that if a given phenomenal property ‘X’ is causally significant in the sense that it expands the

possible states the universe might occupy, it is necessarily causally efficacious because it has been the cause of at least one effect (i.e., perhaps the retrospective phenomenal knowledge that is grounded in our direct self-acquaintance with ‘what it is like’ to be in state ‘X’) that would not have occurred in its absence. Therefore, if X increases the possible states the universe might occupy, then X must necessarily be causally efficacious, and this is because it has, in the simplest possible sense, been the cause of an effect (i.e., our knowledge of a phenomenal event) that would not reasonably have manifested in its absence<sup>56</sup>. Subsequently, either our knowledge of X is rendered explanatorily suspect as we embrace an epiphenomenalism that cannot account for the connection between our knowledge of X and our being in state X, or we hold that X has exerted some causal influence simply by having existed, and this causal influence is manifest in our awareness, or knowledge, of X in such a way that the denial of X’s causal efficacy invariably self-stultifies because every reference to X expands the possible states the universe may occupy and thereby reaffirms its influence. At this stage, even if we maintain that our beliefs pertaining to phenomenal properties are derived from a process of self-acquaintance, being directly self-acquainted with ‘what it is like’ to exist during a given phenomenal event still confers some phenomenal knowledge that expands the possible states the universe might occupy in a direction that reasonably would not have manifested if phenomenality were entirely absent (or truly epiphenomenal), and thereby phenomenality is still causally efficacious because it has still influenced the direction in which the universe flows.

---

<sup>56</sup> If one argues against this, one must return to the problematic inference that, potentially, our awareness, or knowledge, of phenomenality need not be grounded in anything phenomenal at all, and this may very well lead to the contention that phenomenality had no causal significance precisely because perhaps phenomenality did not exist to begin with. However, as iterated in the previous footnote, this is a controversial claim and arguably leads to deeper issues surrounding precisely why (and how) we should seem to be aware of phenomenal properties if they in fact do not exist. It seems to me that, when framed in this way, the most parsimonious conclusion is that phenomenal properties are themselves the cause of our awareness of phenomenal properties, for the alternative leads us to some form of illusionism that introduces far more problems than it solves (see the previous footnotes for a survey of these problems).

Therefore, I hold that it is not unreasonable to posit that phenomenal properties *are* causally efficacious and maintain that this causal efficacy manifests as a necessary quality of both macro-level human subjects *and*, because we have compelling reasons to believe that constitutive reductionism is true and strong emergence, Cosmopsychism and non-constitutive forms of Panpsychism are flawed, also bottom-level microphysical entities<sup>57</sup>. As a result, I argue that we must expand our concept of causation so as find a place for phenomenal causation within our metaphysics, and in so doing must also provide an account of how the causal significance/efficacy of phenomenal properties may be reconciled with our scientific worldview without inviting a commitment to any form of supernaturalism. In the following section, I move to establish phenomenal properties as powerful qualities, in the hope that I can lay the foundation from which to provide such an account.

### 2.1.2 Expanding the concept of causation

Imagine a chessboard replete with various properties and governed by various rules. Now, imagine that these properties and rules are exclusively derived from a physicalist interpretation of the standard model of contemporary physics, in which only spatial-temporal properties and *narrowly* physical causes and effects are said to exist. Given these conditions, although there is a vast quantity of possible outcomes that may arise within the confines of the board, the possible states that the board may occupy remain limited in such a way that

---

<sup>57</sup> See section 1.2.2 for a full articulation of precisely why these are properties of bottom-level microphysical entities and not exclusively emergent properties of higher-order human conscious experience. In brief, I take it that if causal significance/efficacy is an emergent quality of phenomenological properties, we should expect to see a ‘display of some manifestation of the [emergent] presence’ (Papinaeu 2001, p. 31), in a manner that would differentiate these new causal powers from the causal powers entailed by bottom-level micro-phenomenal entities. However, as the same physical laws apply ubiquitously and we see no evidence of new causal powers, we may reasonably posit that if phenomenal properties are causally significant, this causal significance manifests at the bottom-level.

leaves no room for the mental causation that is necessary for the board to behave in a manner that would render it consistent with the nature of the game of chess as human subjects understand it. Similarly, a worldview in which causally significant phenomenal properties are excluded results in a worldview that might be likened to a fragmentary mirror that captures only part of the forces at play but fails to reflect the phenomenal properties which simply must be integrated to offer an account that is consistent with reality as human subjects understand it. Therefore, given the strength of the self-stultifying objection, and the weight of the issues that arise upon extricating phenomenal properties and mental causation from our worldview, I argue that our metaphysical framework must be predicated upon the following maxim:

*Our current epistemic predicament must be ontologically possible, and our current epistemic predicament ought to lead us to an ontological framework that reasonably allows for the possibility of our current epistemic predicament.*

In this sense, the metaphysical constituents must reasonably allow for the possibility of a state of reality in which we can engage in epistemological concerns regarding metaphysical constituents, and our epistemological concerns ought to lead us to a metaphysical framework that is capable of accounting for the possible state in which we raise such epistemological concerns. If we cannot coherently conceive of how our metaphysics necessitates the ontological possibility of this state, we should revise our metaphysics. In this case, a concept of causation that exclusively involves narrowly physical powers seems to lead us to an incomplete worldview that causes us to self-stultify upon denying the reality of non-narrowly physical powers, and therefore we ought to posit that our concept of causation requires expansion on the grounds that if the narrowly physical conception of causal closure were true

then the current epistemic predicament arising from both the complexities of the ‘hard’ problem and the nature of epiphenomenal phenomenal properties would either be highly improbable or absolutely ontologically impossible. I hold therefore that either a narrow physicalist metaphysic results in the impossibility of a predicament concerning the epiphenomenal nature of existent non-physical properties (by virtue of this predicament being caused by properties that directly contravene the ontological commitments of physicalism), or narrow physicalism has (somehow) resulted in a state of affairs in which we are highly confused about the nature of causation as a result of the mistaken belief that we hold phenomenal properties. I suggest that this latter scenario is highly unlikely given that it is not at all clear how (or why) narrowly physical properties would ever produce a deeply mistaken belief of this kind, and so I suggest that it is more likely (and parsimonious) that we are not mistaken about the nature of our consciousness and causation than it is that we are deeply mistaken about that which we are most intimately aware of. I take this line not least because we seem to self-contradict upon attempting to deny the reality of our phenomenal consciousness (or its causal efficacy), but also because once we begin to query our knowledge of that which we know most intimately, how do we then justify not applying this same scepticism to those phenomena which we know least intimately, and thereby how do we avoid the collapse of the entirety of our epistemic framework? I note that whilst this particular argument might not be substantive, I suggest that those who argue for the illusory nature of phenomenality/phenomenal causation (or the collapse of the epistemic problems associated with these phenomenon) simply must address the seeming unlikelihood of a state of affairs in which our current epistemic and ontological predicament has occurred as a result of a referent that neither exists nor occupies the necessary ontology to reasonably be conceivably derived from any (narrowly physical) referent that does exist.

As espoused in the previous chapter, the advantage of Russellian panpsychism is that we are purportedly capable of remaining consistent with the aforementioned maxim, whilst overcoming the otherwise fractured worldview entailed by narrowly physical accounts of causation, by reconciling phenomenal properties with a naturalised account of causation in a manner that maintains the causal efficacy of phenomenal properties without making any discernible difference to the observations recorded in the natural world. Standard accounts of Russellian panpsychism maintain that this is achieved by positing phenomenal properties as a categorical basis of the dispositional properties we observe in the natural world, such that the core of a standard panpsychist form of Russellian monism is: ‘(i) physical properties involve categorical properties, and (ii) consciousness is explained in terms of those categorical properties’ (Goff 2019, p. 102). Therefore, in so far as categorical properties are construed as realising causal powers without being exhausted by a description of their causal role, the Russellian panpsychist may coherently be said to embrace an expanded form of *broadly* physical causal closure that can incorporate a causal role for categorical phenomenal properties into our concept of *broadly* physical causation (Chalmers 2015, Goff 2017/9).

The problem with this is that whilst Goff (2019) defines categorical properties as properties that may ‘involve, or realise, causal powers [whilst being] more-than-merely-dispositional nature’ (p. 100), it is not immediately clear precisely how such categorical properties may possess the necessary causal powers to ground the dispositions we observe in the natural world, and further still it is less clear precisely how we may indubitably demarcate that these categorical properties are *phenomenally* causally relevant in such a manner that the causal significance of phenomenal experience is rendered explicit.

These issues are elucidated in detail by Howell (2014), who suggests that any form of Russellian panpsychism predicated upon placing phenomenal properties as a categorical ground for the dispositions we observe in physics invariably faces substantial explanatory issues. According to Howell (2014), whilst phenomenal properties may coherently act as the categorical ground for dispositional properties, ‘every physical property still has a sufficient physical cause’ (p. 31), and therefore (narrowly) physical ‘dispositional properties are causally sufficient for the appearance of dispositional properties’ (ibid, p. 31). This is coherent in so far as a Russellian panpsychism that construes phenomenal properties as categorical does *not* necessarily commit us to the thesis that dispositional properties necessarily involve anything more than narrowly physical properties, and this is what Howell (2014) believes works against Russellian panpsychism (p. 32). The argument follows that whilst we may grant that phenomenal categorical grounds are causally relevant in so far as there must be *something* that confers dispositions and causal powers, these categorical grounds are complex and thereby must ground both phenomenal *and* causal resemblance relations.

Howell (2014) is therefore positing that as multiple dispositions arise because of a categorical ground, the phenomenal aspects of a categorical ground may not necessarily be responsible for the physical dispositions that manifest as physical, causal resemblance relations. Howell justifies the move to highlight the separability of causal resemblance relations from their phenomenal counterparts by appealing to a modal argument in which the phenomenal experience of ‘green’ shares a phenomenal resemblance in world 1 and world 2, but grounds distinct causal powers in each such world. Such that, for example, the experience of green grounds negative charge in W1, but the same experience grounds negative spin in W2. Howell maintains that as our categorical ground is phenomenal, the

phenomenal resemblance relation must conceivably hold across all worlds, but a phenomenal categorical ground can conceivably produce distinct causal powers in distinct worlds replete with distinct physical laws. Further, as this form of Russellian panpsychism commits us to phenomenal properties as the categorical ground of dispositions, these phenomenal properties must be responsible for a phenomenal resemblance relation that holds universally *and* equally responsible for a causal resemblance relation that holds contingently. As such, we may separate aspects of our phenomenal, categorical ground into those aspects that ground universal phenomenal resemblance relations and those aspects that ground contingent causal, physical resemblance relations.

With this established, Howell (2014) advances his central argument:

1. *There are two distinct and separable aspects of Russellian Panpsychist (RP) properties, those that ground phenomenal resemblance relations and those that ground resemblances between causal profiles;*
2. *All physical events have sufficient causes in virtue of those aspects that ground resemblances between the causal profiles of RP properties.*
3. *Therefore, the aspects of RP properties that ground phenomenal resemblances make no unique causal contribution to the physical world. (p. 33)*

This argument concludes that the phenomenal aspects of categorical grounds are resigned to the epiphenomenalism Russellian panpsychism was attempting to avoid, on the grounds that the (narrowly) physical aspects of categorical properties act as sufficient grounds for physical events. In this sense, Howell is arguing that whilst the phenomenal aspects of categorical properties may ground phenomenal resemblance relations, these phenomenal aspects remain causally inert in relation to the physical because physical events are



sufficiently causally explained by the aspects of categorical properties that are non-phenomenal. Therefore, whilst Russellian Panpsychism is motivated by an attempt to show that categorical properties hold causal powers in virtue of their phenomenality, ‘phenomenal natures once again end up making no unique contribution to the physical world’ (Howell 2014, p. 33).

I posit that this problem *is* intractable and has occurred because the narrative presented thus far (the narrative espoused by Goff and others) has focused upon the possibility that the dispositional properties analysed by physics are phenomenally grounded in the categorical properties of the physical. Whilst *prima facie* this seems to present a coherent worldview by avoiding the issues involved in positing phenomenal properties as either purely dispositional properties<sup>58</sup>, or as purely causally inert categorical properties, it remains difficult, using this model, to fully overcome Howell’s objections. Arguably, Howell’s objections remain steadfast because viewing phenomenal properties as a categorical ground for physical dispositions leaves us with, what Martin (1993) references as, a ‘two-sided coin’ (p. 184) view of properties, such that phenomenal properties present dual aspects: one aspect that grounds phenomenal resemblance relations *and* another that grounds physical resemblance relations. The problem with this, as Howell (2014) rightly points out, is that we must clarify whether these two aspects are related contingently or necessarily. Howell seems to rightly maintain that if one aspect grounds a phenomenal resemblance relation necessarily, but the other aspects grounds causal resemblance relations contingently, then it is unclear precisely how we may substantiate the notion that phenomenal properties ground

---

<sup>58</sup> For dispositional properties are known only by virtue of their relations with other dispositional properties, and it seems exceedingly difficult to account for the qualitative aspects of subjective phenomenal experience within a purely relational model of properties.

causal dispositions because, if they did, we should expect to see the same phenomenal resemblance relation causing the same causal resemblance relations in all possible worlds.

Framed as such, the most immediate way out of Howell's conundrum is to posit that the relation between these two aspects *is* necessary, and therefore if a phenomenal resemblance relation "W" holds, a causal resemblance relation "X" also holds. The most elegant means of achieving this is to endorse a version of the powerful qualities view advanced by Heil (2013), in which 'a property's dispositionality and qualitativity are not aspects or properties of the property, they are rather the property itself, differently considered' (p. 212). Heil and Robb (2003) use the analogy of a square to compound the possibility of powerful qualities:

*'Consider an ordinary quality: being square. This quality might appear to be a clear example of a categorical property. But note: in virtue of being square, an object is disposed to pass smoothly through holes of certain shapes (and not through others), disposed to reflect light in a particular way, disposed to make an indentation of a particular kind in a lump of clay. Being square, then, appears through and through dispositional.'* (p. 185-186)

When combined with Russellian micropsychism, this view entails that there is fundamentally a category of phenomenal-physical properties that ground all the powerful qualities we observe in physics, and this property presents two modes of presentation: qualitative modes of presentation manifesting as phenomenal character when viewed from the inside and dispositional modes of presentation manifesting as structural relations when viewed from the outside. Consequently, categorical and dispositional properties are two modes of

presentation<sup>59</sup> for this phenomenal property that is *both* categorical and dispositional, and therefore categorical and dispositional properties are identical. Construed as such, there is one fundamental type of physical property that just *is* phenomenally experiential<sup>60</sup>, and ‘talk of a dispositional property is just a way of characterizing the way the law governs that [one property] in relation to other properties’ (Howell 2014, p. 26)<sup>61</sup>. By employing a view of this kind, we may immediately demarcate numerous qualities of this property, for example we may analyse its spatiotemporal dispositions using the natural sciences we currently employ, and, by virtue of the relations being identical to a given phenomenal state, we may reasonably make at least some inferences pertaining to a structure’s intrinsic, phenomenal nature. Indeed, the advantage of this theory of properties is that it accounts for phenomenality as a causal ground for physical resemblance relations whilst remaining entirely naturalised. This is because the spatiotemporal relations analysed by the natural sciences are simply imbued with a phenomenal quality that *is* the dispositions manifested within spatiotemporal relations, and therefore a further benefit of endorsing powerful qualities is that the natural sciences may themselves provide a means from which we may make inferences pertaining to the nature of such powerful (phenomenal) qualities. As a result, the thesis I am advancing, in line with advocates of powerful qualities and Russellian Micropsychism, is that categorical and dispositional properties are themselves nothing beyond two modes of presentation for a singular, unitary ontologically fundamental phenomenal-physical property, and this property may be presented either spatially from the outside or experientially from the inside<sup>62</sup>.

---

<sup>59</sup> Much the same way that even though Hesperus is Phosphorus, our modes of presentation cause us to denote this star using two distinct concepts dependent upon the time of day.

<sup>60</sup> One problem with this account appears to be that it is difficult to conceive of disparate qualities/dispositions arising from this singular property. I address this issue in detail within chapters 3 and 4.

<sup>61</sup> What precisely this means shall be extrapolated in the subsequent chapter.

<sup>62</sup> See Jacobs (2011) and Schneider (2017) for similar articulations of the thesis that turns on the notion that phenomenal properties are powerful qualities.

In so doing, I concede to the inference that the relation between categorical and dispositional properties is necessary, and thereby simultaneously entirely avoid Howell's issues by also conceding that it is inconceivable that an identical instantiation of phenomenality would manifest distinct dispositions in various possible worlds<sup>63</sup>. Whilst *prima facie* the necessity of this inference does not seem to present an issue, on closer inspection it requires a full explication of a naturalised account of phenomenal properties that is capable of explaining precisely what dispositions may arise by virtue of being phenomenally experiential and also an account of precisely why it is impossible for the phenomenal property "X" in W1 to ground the disposition "Y" whilst the exact same phenomenal property in W2 grounds the disposition "Z"<sup>64</sup>. Both such accounts will be evident upon a deeper explication of what precisely phenomenal properties are and will therefore be provided within the subsequent sections. In this vein, I end this section by noting that of the possible options available to us in the wake of the self-stultification principle, the concept of powerful qualities allows us to simultaneously 1) explain phenomenal causation, 2) explain the relationship between dispositions and phenomenal properties, and 3) achieve all of this without contradicting our otherwise robust naturalised worldview. In light of this, I posit that we ought to be absolutely unconcerned with the question of whether phenomenal properties play a causal role, and instead ought to concern ourselves exclusively with delineating precisely what

---

<sup>63</sup> As shall be made clear in the subsequent sections of this chapter (and parts of chapter 3), I take this to be inconceivable because an instantiation of phenomenal content is identical to the given structure that this instantiation of phenomenality *is* (and the nature of these structures is defined by the natural laws at play in each possible world). Subsequently, it is inconceivable that the exact same instantiation of phenomenality should ground distinct dispositions in different possible worlds because it is entirely reasonable to surmise that identical structures would not manifest in distinct possible worlds (and subsequently the instantiations of phenomenality in each such world could not reasonably be described as identical). In this respect, as it is inconceivable that truly identical instantiations of physical structure should occur in possible worlds with distinct natural laws, it is equally inconceivable that identical instantiations of phenomenality should occur in such worlds, and thereby, by positing the identity of an instantiation of phenomenality as, in the relevant sense, one and the same as a given structure, we seem to avoid Howell's concerns.

<sup>64</sup> To reiterate the previous footnote, I take it that this problem will be addressed as soon as the relationship between phenomenality and structure is rendered transparent in section 2.3, because for an instantiation to be identical in another possible world, its structure must also be identical, and conceiving of how identical structures might manifest amidst distinct natural laws is extraordinarily difficult, if not impossible.

role phenomenal properties may play<sup>65</sup>. By means of achieving this, I continue with my account of the nature of phenomenal properties by presenting the case for construing all such properties as ‘experiences-for’ a subject.

## **2.2 “Experiences-for”: On subjects of experience and the impossibility of unexperienced instantiations of phenomenality.**

At this moment in space-time, if there is something it is like for you to read this thesis, we might justifiably hold that you have phenomenal properties. From this, we might extrapolate that you are a subject of experience by virtue of being in this phenomenal state. When discussing phenomenal properties, I will, as I have done throughout, ascribe them to any entity ‘S’ in which a subjective experience of ‘what it is like’ to *be* ‘S’ manifests (in this sense, my concept of such properties is largely Nagelian 1974 in nature). It is of import to note that I am explicit that these are properties *of* entities replete with subjectivity and *not* simply properties of entities broadly construed. In this sense, it seems to be a matter of necessity that the qualities of phenomenally conscious states *are* phenomenological properties if phenomenological properties are taken to be subjective conscious properties in and of themselves. This definition is, for the most part, uncontroversial and endorsed by the majority of contemporary philosophers of mind, with most holding that upon instantiating phenomenal properties you simply *are* a particular subject of experience (see Speaks 2015, Chalmers 2013 amongst others) in such a way that that upon discussing the phenomenal property of ‘redness’, for example, we are *not* inferring that entity “S” undergoes an experience of phenomenal redness that belongs to a more ontologically fundamental subject.

---

<sup>65</sup> See chapter 3 for a full account of the sort of powerful qualities I am presented and an account of how imbuing powerful phenomenal-physical phenomenal into our worldview might be reconciled with our scientific models of reality whilst also making an explicit difference to our understanding of the natural world and its workings.

Instead, by virtue of having a phenomenal experience, or instantiating phenomenal properties, the entity necessarily simply *is* a subject of experience to the extent that if this instantiation of phenomenality is posted as fundamental, so too is the subject that realises it.

Whilst this definition, and its associated commitments, might appear to be intuitive, there are some philosophers who reject the notion that instantiations of phenomenality necessarily entail subjectivity (see Coleman 2014/18), on the grounds that we can demarcate phenomenal qualities (qualities of subjective experience that do not necessarily require a subject to exist) from phenomenal properties (subjective experience in and of itself). To disambiguate this, let us consider the perceptual experience involved in the observation of this page. In this case, the perceiver observes a white background with various black markings. The issue at hand is to delineate whether the white-ness/black-ness phenomenal qualities of this experience are phenomenal properties (properties of subjective conscious experience that may not exist non-phenomenally) or white/black phenomenal qualities of the page itself. *Prima facie*, it does appear at least conceivable that we are simultaneously referring to, on the one hand, a phenomenal colour property contained in the page and, on the other hand, a property of experience. In this sense, we immediately reveal two conceptions of the term ‘phenomenal property’: 1) the property involved in this experience of white-ness is an essential feature of our phenomenology – it cannot coherently exist devoid of subjective ‘what it is like’ experience and is thereby identical to a type of phenomenal experience, 2) the quality involved in the experience of white-ness is a feature of the page that might persist even in the absence of a phenomenal experience. Construed as such, Coleman’s motives become slightly clearer, for when we are discussing the phenomenal whiteness of our experience of this page, we are faced with something of a conceptual hurdle. As Montague (2014) elucidates, we are referencing, on the one hand, ‘a purely phenomenological property of experience, a property

that is not, in fact, a property of non-phenomenological things in the world', but there is also a 'fundamental *sense*' in which we attribute phenomenal whiteness to the page (p. 45). In this latter case, the whiteness (what Montague 2014 might refer to as a 'whiteness-as-seen') is a property that is experienced as 'a mind-independent' property of the page, 'not as a property of the experience' (Montague 2014, p. 45). In this respect, there really does seem to be two senses employed to refer to the concept of phenomenal properties. On one hand, these might refer to 'as seen' qualities of non-phenomenological objects, and on the other, they might refer to a property that is an intrinsic feature of phenomenological experience and is thereby decidedly (and necessarily) *not* a property of non-phenomenological, mind-independent objects.

In what follows, I argue that this perceived disparity is reducible to a conceptual misunderstanding that arises amidst our attempts to delineate the essential features of our experience. In this case, I show that referring to phenomenal qualities as qualities of non-phenomenal objects redefines phenomenal qualities to such an extent that they are rendered largely ambiguous. For the most part, this critique turns on the contention that as phenomenal qualities refer to qualities of experience, it makes no sense whatsoever to erroneously employ this concept to position qualities of experience as features of the world that exist devoid of an experiential locus in which these qualities are realised. Explicated as such, I maintain that the conceptual tools we employ in the discussion of subjective conscious experience must be disambiguated, and in so doing we must address that by a matter of natural fact both phenomenal properties and phenomenal qualities (broadly construed) are one and the same in so far as both refer to subjective conscious experience. Subsequently, I hold that whilst it may prove useful to employ two concepts to disambiguate subjective conscious experience from the qualities of subjective conscious experience, it seems inevitable that, ultimately, we are

incapable of isolating phenomenal properties from a fully realised concept of phenomenal qualities (or vice versa) without rendering our concepts unintelligible.

In order to ground this point, I recognise two senses for the term ‘phenomenal property’: 1) a ‘minimal’, restricted sense in which our *prima facie* intuitions lead us to construe them as mere qualities of subjective experience and thereby conceive of them as properties exclusively employed to delineate between types of qualitative content, such as, for example, the specific qualities of ‘whiteness’ or ‘blackness’ that might conceivably exist even in the absence of subjective experience, and 2) a ‘broad’ sense in which we employ rational reflection to conceive of them in totality and thereby reveal that they cannot be anything but phenomenological properties, to the extent that when referencing any instantiation of phenomenal experience we are necessarily referring to subjects of experience, and their subjective experiential states, in and of themselves<sup>66</sup>. I maintain that the first ‘minimal’ elucidation is the cause of the move to position phenomenal qualities as non-phenomenological qualities of objects, for it is, as I shall argue, this confusion in our understanding of phenomenal properties that has enabled the move to separate the qualities of phenomenal experience from the subjective experiential fields that ground said qualities.

As intimated, Coleman (2014/17) stands as one of the few philosophers of mind who maintains that phenomenal qualities are necessarily ‘minimal’ and argues that much in the same way that unconscious thoughts may ostensibly exist, there may be phenomenal qualities that exist devoid of an experiential subject. With this as grounding, Coleman (2014/17) maintains that 1) we may disambiguate phenomenal qualities from subjects of experience,

---

<sup>66</sup> I have borrowed the implication that there may be these two senses in which we conceive of concepts from Chalmers (2010), although I note that Chalmers did not have phenomenal properties in mind.



and 2) upon doing so, we may posit a panqualityism capable of providing a means of solving the combination problem afflicting standard forms of panpsychism. As the coherence of this latter claim rests exclusively upon the coherence of the former, I take it that upon addressing the impossibility of conceiving of unexperienced phenomenal qualities, I also address the impossibility of employing a fundamentally experience-less and subject-less form of panqualityism as a means of coherently accounting for the singular unified conscious experience that is at the core of the standard combination problem. As such, I devote this section to addressing the impossibility of unexperienced phenomenal qualities (and therefore also phenomenal properties), and in so doing frame an account of phenomenal properties that is necessarily ‘experience-for’ in the sense that any instantiation of phenomenality is necessarily an instantiation of a subject of experience.

Coleman (2014/17) advances his case for the possibility of unexperienced qualia by championing Rosenthal’s (1991) critique of the Nagelian (1974) position that upon instantiating a phenomenal property, for example, the experience of ‘redness’, a singular conscious subject who experiences this ‘redness’ must also be instantiated. For the most part, Nagel’s position falls in line with the zeitgeist of organised philosophical thought, and this is because if a phenomenal property is an experience of ‘what it is like’ to undergo, for example, ‘redness’, we can only make sense of such a property in so far as this is a ‘what it is like’ experience that manifests as a singular point of view *for* a ‘unified and bounded’ subject of experience<sup>67</sup>. As such, a quality of experience is said to exist only in so far as there exists a subject to whom it may be ascribed. Consequently, for the most part, contemporary philosophy of mind maintains that phenomenal qualities entail a singular, subjective point of

---

<sup>67</sup> That is, an experiential field that is decidedly unified such that it presents to one singular point-of-view, or ‘subject’, and bound to this ‘subject’ “S” in a manner that is not shared by all those subjects not S.

view manifesting as a phenomenally conscious subject of experience for whom the quality may be realised<sup>68</sup>. In this sense, when we are discussing fundamental experiential qualities, we are discussing a subjective quality of experience, and are thereby, at foundation, discussing fundamental subjects of experience.

As a result, the standard interpretation of phenomenal properties leaves us with an entangled view of phenomenality, such that upon discussing phenomenal properties we are referencing the qualities of a given experience and, by virtue of their experiential nature, necessarily also referencing an experiencer to whom the qualities may be ascribed. I take it that this entangled view is necessary, for it seems the essence of a phenomenal quality is the manifestation of a feeling that necessarily consists in being felt by an experiential subject, as Reid (1855) elucidates:

*This sensation can be nothing else than it is felt to be. Its very essence consists in being felt; and when it is not felt, it is not. There is no difference between the sensation and the feeling of it; they are one and the same thing. It is for this reason...that in sensation there is no object distinct from that act of the mind by which it is felt; and this holds true with regard to all sensations.* (p. 141)

Contrarily, Rosenthal (1991) and Coleman (2014/17) maintain that this entangled view of phenomenality leaves us with an ‘impossible conceptual knot’ (Coleman 2014, p. 68) that must be unravelled into two opposing threads: ‘consciousness on the one hand, and the

---

<sup>68</sup> As articulated in the previous chapter, to truly make sense of the type of subjective experiences we associate with phenomenal properties, we must maintain that subjective experience is necessarily understood as a unified and bounded experience from a particular unified and bounded point of view, i.e., an experience ‘from one type of point of view: that of a being like that one having an experience’ (Nagel 1979, p.188).

qualities one is conscious of, on the other' (ibid, p. 69). This 'dual-aspect' view, then, defends the distinction between phenomenal properties (the property of subjective consciousness in and of itself) and phenomenal qualities (qualities of subjective consciousness), with the caveat that both Rosenthal (1991) and Coleman (2014/17) view consciousness as a higher-order state of cognitive awareness and view phenomenal qualities as the sensory qualities that are represented during a particular higher-order perceptual state. Therefore, if an entity does not possess the higher-order brain state necessary for awareness, an entity may be undergoing a sensory quality such as redness without being conscious of their doing so<sup>69</sup>. Rosenthal and Coleman maintain that if this is possible it is not incoherent to posit unexperienced phenomenal qualities, for a given instantiation of phenomenality, for example 'redness', may manifest as the consciously experiential quality of redness to certain higher-order subjects whilst manifesting as a quality devoid of consciously experienced qualitative features to others. As Coleman (2014) puts it, if we can conceive of phenomenal qualities instantiating consciously experienced qualitative features in some entities and not others, and if we are realists about colours to the extent that the phenomenal quality of redness is not necessarily a quality *of* an experiential subject but is instead a quality of a given entity that exists independently of perceivers, then we ought to be committed to the thesis that phenomenal qualities may endure when unexperienced (p. 29)<sup>70</sup>.

The problem Coleman faces, however, is that even upon embracing an account of consciousness in which it is positioned as nothing over and above a physical, functional

---

<sup>69</sup> As Coleman (2014) notes, this allows Rosenthal to explain cases of blindsight, in which individuals respond to sensory stimuli that they do not consciously perceive (see p. 69-70).

<sup>70</sup> The core of Coleman's contention here rests upon the coherence of construing consciousness as a higher-order relation, for in an attempt to substantiate this move to establish the conceivability of unexperienced qualia Coleman employs an account of consciousness that is entirely structural and relational in nature and rests exclusively upon the concept of awareness, such that 'consciousness is simply that relation whereby qualities are brought into awareness for their bearer' (p. 31).

mechanism by which awareness (and thereby consciousness) is achieved by virtue of interacting with these mind-independent phenomenal qualities<sup>71</sup>, he must still account for precisely how the physical mechanism that accounts for awareness ever achieved the sort of states we associate with subjective experience. How, for example, did the purely physical mechanism ever achieve awareness of undergoing an experience of ‘what it is like’ to witness the blueness of the sky? Answering this question is what led Coleman to endorse his panqualityist form of neutral monism that embraces the possibility of qualia externalism, and thereby grants 1) that ‘all matter possesses [irreducible] qualia, which need not be experienced in order to exist’ (p. 29), and 2) certain higher-order states of awareness may represent the phenomenal qualities of matter. Therefore, it is matter itself that is responsible for ‘embedding [the] qualia-carrying sensory states’ (Coleman 2014, p. 32) that account for the physical mechanism’s awareness of subjective experiences<sup>72</sup>, and as a result when external phenomenal qualities interact with the awareness imbued by a sufficient relational mechanism, Coleman (2014) ostensibly ‘capture[s] that aspect of subjectivity which is the sheer awareness of qualities—a.k.a. consciousness’ (p. 33).

Whilst there are various issues with a functionalist qualia externalism of this kind, I do not deem it necessary to address them all here<sup>73</sup>. Instead, I focus with rigidity upon Coleman’s (2014) contention that phenomenal qualities ‘need not be experienced in order to exist’ (p.

---

<sup>71</sup> I note that if we were to allow for the possibility that a purely physical mechanism of this kind is possible, this move substantially distances Coleman from the standard panpsychist interpretation of consciousness as a minimal, phenomenal subjectivity that manifests as ‘what it is like’ experiences, such that “an organism has conscious mental states if and only if there is something that it is like to be that organism—something it is like for the organism” (Nagel 1974, p. 436). Construed as such, the disparity between identifying consciousness as a first-order property and second order property is evident. The former posits that consciousness may exist minimally and fundamentally; the latter maintains that consciousness only occurs upon sufficient structural relations emerge to force ‘the material world back upon itself, to apprehend its own character’ (Coleman 2014, p. 31).

<sup>72</sup> In this sense, Coleman’s thesis is not wholly distinct from the theses of other qualia externalists, such as Tye (2000).

<sup>73</sup> See Hibbert (2018, p. 90) for an articulation of the issues associated with similar accounts.

29) and his subsequent inference that this is justified by virtue of the fact that, to Coleman's knowledge, 'no panpsychist has provided a serious argument for the claim that qualia (phenomenal qualities) cannot exist unexperienced' (ibid, p. 30). I take it that such an argument might be provided upon highlighting the tenuous nature of the contention that the concept of an unexperienced phenomenal quality is a concept that holds without contradiction. This is something Coleman and his contemporaries arguably have great difficulty with, especially as we can arguably only coherently conceive of an instantiation of phenomenality existing as an experience of some type, and we can only conceive of an experiential type existing *if* there is an experiential locus in which the type manifests as an experience *for* a singular experiencer. I take this to be largely uncontroversial on the grounds that phenomenality, and the qualities of phenomenality, seem to necessarily entail a subjective aspect that manifests as a unique, subject-specific experience for a specific experiencer, and therefore it seems that it makes very little sense to suggest that phenomenal qualities might exist as antecedents to a subject capable of realising them, for in order to exist these qualities arguably require an experiencer – a subject – 'for whom there is something it is like for that experience to exist at all' (Guillot 2016, p. 27).

In a not dissimilar vein, Zahavi and Kriegel (2015) argue that to deny the subjectivity of phenomenal qualities is to misconstrue phenomenal qualities entirely, for it is extraordinarily difficult to construe phenomenality manifesting in anything than a 'distinctly first-personal way', and as a result 'to deny that such a feature is present in our experiential life, to deny the for-me-ness or mineness of experience, is to fail to recognise the very subjectivity of experience' (p. 38). To say, therefore, that the phenomenal quality of 'redness' (that is, the 'what it is like' experience of 'redness') is contained in mind-independent entities, and thereby may exist devoid of subjectivity and unexperienced, is therefore entirely incoherent

on the grounds that, in lieu of an argument to the contrary, we cannot conceive of how an experience may be grounded in anything other than an experiencer. This is because, in order for the concept of phenomenal experience to coherently satisfy what we commonly mean by an instantiation of phenomenality (or a phenomenal quality), it must be grounded in a sufficiently experientially bounded and subjectively unified locus in which the ‘what it is like’ experience instantiates a spike of subjective ‘feeling’ that is ‘felt’ by a specific entity “S” and manifests as ‘something it like’ to be “S” and not all those entities “~S”, and if there is *not* ‘something it is like’ for ‘S’ in this sense, it is difficult to see how the quality can be described as ‘phenomenal’. Without positing an experientially bounded locus (or subject) of this type, it becomes incredibly difficult to conceive of how we may construe phenomenal experiences at all, for it seems if subjective, phenomenal experiences are themselves *not* beholden to a particular subject, they are no longer subjective experiences at all. In this respect, I argue that we cannot hope to disentwine subjectivity, or the subjective ‘feel’ of what it is like to undergo an experience, from the phenomenal qualities themselves because it is the very subjective character of an experience that realises these qualities as qualities replete with phenomenal content. And, subsequently, it is impossible to disentwine subjects from phenomenal qualities without either rendering phenomenal qualities inconceivable or redefining them to denote something other than subjective conscious experiences.

I suggest therefore that, by necessity, a feeling is nothing if not felt, just as a subjective quality of experience is nothing if not experienced, and so by attesting that the qualities of ‘what it is like’ experiences may exist unexperienced, Coleman seems to invite the contradiction entailed in holding that a subjectively experiential quality ‘S’ exists as a subjectless, unexperienced quality ‘~S’. Construed as such, unexperienced phenomenal

qualities appear inconceivable (and thereby impossible) on the grounds that entertaining their possibility reveals a contradiction manifesting within the notion of a quality that simultaneously instantiates subjective ‘what it is like’ experiences *and* the absence of a subject in which subjective experiences may manifest. In this sense, I maintain that Coleman cannot realistically avoid falling foul of the contention that if a phenomenal quality (Q) is instantiated, then necessarily a phenomenological property (P) is instantiated because our understanding of Q is such that (Q and  $\sim$ P) is only possible if Q denotes something other than a particular experience of ‘what it is like’ *for* a subject. It seems extraordinarily difficult for Coleman to coherently avoid the necessity of this relationship, however. Especially as phenomenal qualities seemingly must manifest as ‘unified and bounded’ content experienced from a particular subjective point of view in order to coherently capture the qualities of subjective phenomenal experience and thereby satisfy our concept of what a phenomenal quality actually *is*.

Construed as such, it seems all Coleman may coherently argue is that certain structures may act as *potential* carriers of proto-phenomenal qualities, which may transition into fully fledged phenomenal qualities upon interacting in the right way with a subject of experience, but that these ‘proto-phenomenal qualities<sup>74</sup>’ (or mere intentional properties) are themselves *not* pillars of phenomenal experience and are thereby *not* phenomenal qualities properly construed. Interestingly, by persistently employing the concept of *unexperienced* phenomenal qualities, Coleman does seem to opt for redefining phenomenal qualities to denote something other than subjective phenomenal qualities. I

---

<sup>74</sup> I note that proto-phenomenal qualities are distinct from proto-phenomenal properties in so far as proto-phenomenal properties are not phenomenal properties but might (somehow) realise phenomenal qualities and subjectivity, whereas proto-phenomenal qualities are not phenomenal qualities but might (somehow) realise only the qualitative content of phenomenal states.

note, however, that if this is Coleman's aim, he should cease to refer to these non-phenomenological properties as *phenomenal* qualities and should instead simply refer to them in one of the two senses described above, especially as doing so would avoid the conceptual ambiguity entailed in the usage of phenomenal qualities to seemingly simultaneously refer to instances of non-subjective experience *and* subjective experiences. I note, however, that Coleman might very well be reticent to do this on the grounds that it is only those *phenomenal* aspects of the qualities he references that can be employed to solve the problem of how we transition from an aware non-phenomenal mechanism, or functional brain state, to a subjective conscious experience. So, describing the qualities Coleman references as what they are – mere intentional properties – might be rejected by Coleman because this provides no explanatory purchase to his problem of accounting for how subjective conscious experience arose in the first place.

In this respect, even on a generous reading of Coleman's account in which we concede to the possibility of conceiving of unexperienced qualia, we might still push for a solution to the problem of how, if phenomenal qualities truly exist devoid of a subject of experience (as Coleman ardently maintains), the subject-less, experience-less relational mechanism that accounts for awareness stands in the sort of relation to subject-less, experience-less phenomenal qualities necessary to produce a *subjective experience* that exists as anything above and beyond an illusion? Surely, if these phenomenal qualities are fundamentally unexperienced and subject-less, we cannot derive true subjective experience from standing in relations to these qualities. And, surely, if the mechanism itself is simply an experience-less, subject-less functional machine that stands in some relation to the phenomenal qualities of entities, we cannot expect to derive subjective



experience from the functioning of this mechanism either. So, it is unclear precisely how an experience-less, subject-less functional machine (regardless of complexity) standing in this relation to subject-less, experience-less phenomenal qualities should cause subjective experiences in that which is not in itself already occupying a subjective field of experience. Indeed, maintaining that experience-less phenomenal qualities may themselves interact with an experience-less mechanism, and thereby be the cause of subjective experience, is akin to embracing a strong emergentism by attesting that the right combination of experience-less entities will somehow ‘traverse the magic passage across the experiential/non-experiential divide’ (Strawson 2008, p. 70).

Consequently, in accounting for subjective conscious experience proper, I suggest that subjectivity simply must be imbued within the fundamental constituents themselves and, as a result, we ought to reject the notion of unexperienced qualia (and subsequently also reject the ‘minimal’ sense of phenomenal properties described at the outset of this section). Therefore, because providing an account of phenomenal properties devoid of subjectivity provides no explanatory purchase to the problem of subjective conscious experience, and because we arguably cannot coherently conceive of phenomenal properties without an *experience-for* relation, these properties simply must entail subjective conscious experience in such a way that subjectivity is construed as a necessary and fundamental feature of any truly phenomenal state, and, as a result, by positing a Russellian Micropsychist account of ontologically fundamental phenomenal properties, we necessarily posit bottom-level microphysical entities as fundamental subjects of experience. Subsequently, a phenomenal property is the property of there being ‘something it is like’ *for* a subject of experience (an *experience-for-a-subject*), and so, at core, is a property that denotes both an experience and an experiencer (a subject of

experience). With this established, in order to fully articulate the nature of fundamental phenomenal properties, we must now articulate the nature of fundamental subjects of experience.

### 2.2.1 Subjects, Subjects of experience and Subjectivity

To provide such an account, it seems befitting to first demarcate what precisely is meant by the term ‘subjectivity’, and in so doing demarcate the perceived disparity that exists between subjects (broadly construed) and minimal subjects of experience of the type that might be entailed by bottom-level phenomenal properties. As briefly highlighted in the previous chapter, I take subjectivity to exemplify any instance of content that is ‘unitarily bounded’ to a particular entity ‘S’ in such a manner that this content (Y) is not shared with all those entities not S. In order to fully articulate this point, I reiterate that when discussing the concept of content being unitarily bounded, I am discussing content that is both expressly unitary in such a way that it manifests to a singular unitary whole that is not in itself disparate or fractured (and may therefore be coherently posited as a singular point of view, subject, or entity ‘S’), and bound to this particular unitary subject in such a way that S and only S undergoes this particular manifestation of the content Y, such that the content Y is not split or shared by all those subjects not S. It is precisely this that I take to denote subjectivity, and, with this established, I am now able to begin to discriminate between broadly construed subjects, of the type that both I and the reader of this piece are likely to identify with, and minimal subjects of experience, of the type that may be ascribed to bottom-level entities<sup>75</sup>.

---

<sup>75</sup> Prior to this attempt to discriminate between these two types of subjects, I preface my account by noting that I take it that both such articulations share the property of subjectivity on the grounds that the concept of being a subject is itself only meaningfully communicated if a given subject ‘S’ is unitarily

In all theories of the phenomenal subject, the disparity that lies at the heart of competing conceptual frameworks is whether the experience and the subject are distinct particulars, or whether they are not. This is what it boils down to: the subject is either a recipient of an experience, or the subject is identical with, or a necessary feature of, the experience. In contemporary philosophy, it is common to adopt the former conceptualisation and thereby attest that the subject is ontologically distinct from the experiences it undergoes. This stance identifies the subject with the functional mechanisms necessary to ground self-awareness and thereby argues that this awareness, whatever it may be, is the foundation of the subject on the grounds that experiences are presented to an awareness that is itself self-aware enough to know that there seems to be a relation between the self-aware subject and the experiences this subject undergoes. Further, if this *is* a relation, the subject cannot be ontologically identical with experiences and thereby may conceivably persist in the absence of all experience. In this sense, we might, in line with Guillot (2017), define the distinction between these two ways of conceiving of the subject as follows: on the former, the phenomenal subject is positioned as ‘a quality or feature our experiences have or come with which suggests a monadic predicate, of the form  $F(x)$  (where ‘ $x$ ’ ranges over experiences), whereas on the latter the subjective character would seem instead to assume the form of a relation: a relation  $R1$  of awareness between a subject  $s$  and an experience  $x$  of hers, of the form  $R1(s,x)$ ’ (p. 33). With this distinction now in place, in what follows I address the coherency of this higher-order,

---

bounded in such a way that it is distinct and irreducible to all those subjects  $\sim S$ . This is the case because if a subject  $S$  is not distinct from all those entities  $\sim S$ , then we lose the mechanism by which we may meaningfully ascribe the moniker of ‘subject’ to a given entity, for the concept itself involves a subject-object relation that dissolves, in the vein of Heidegger (1927), as soon as we disallow the reality of subject-object distinctions. I take it therefore that as we *can* meaningfully discuss the concept of subjects, we reveal, at least in some sense, the reality of a subject-object distinction, and therefore I hold that, at foundation, *subjects are necessarily unitarily bounded instances of subjectivity*, and maintain that if this point is contested, the onus lies with the contestor to offer an account of subjects that does not, by necessity, entail a subjectivity of the type I have described.

relational conception of the subject, before advancing an argument designed to weaken the credibility of any account that attempts to entirely disentwine subjects from experiences. From here, I suggest that an account of Russellian Micropsychism might avoid these issues by incorporating subjectivity at the bottom-level, and, with this established, in the subsequent sub-section I delineate the disparity between a macro-level and micro-level conception of subjects and offer an account of the sort of ‘minimal’ subjects that might reasonably be entailed by bottom-level phenomenal properties. Finally, I elucidate the difficulties that one of the most championed definitions of minimal subjects faces, before finalising this section by offering a means of expanding the concept of minimal subjects to avoid some of the issues facing this account (and avoid the issues afflicting relational accounts). To reach this point, I begin by highlighting the explanatory issues that occur upon endorsing a relational account of subjects.

The most immediate difficulty with such accounts, as I outlined in the previous section, rests in attempting to fully and coherently explicate precisely how a subject may be disentwined from an ‘experience of hers’, for by necessity if this experience is truly *hers* such that it cannot be anything other than an experience with a particular subjective character *for* her, then there is no means whatsoever in which the experience may be disentwined from the subject. As Kriegel (2009) attests: ‘a conscious experience’s qualitative character makes it the conscious experience it is, while its subjective character makes it a conscious experience at all’ (p. 101). Therefore because the notion of phenomenal consciousness cannot coherently be separated from the notion of subjective character (i.e., the particular subjective feel of an experience for this particular subject ‘S’ as distinct from the feel of an experience for all those subjects  $\sim$ S), it seems that phenomenal consciousness refers to any instance of

subjective conscious experience, and as such ‘subjective character is simply another name for phenomenal consciousness, considered generically and under the most neutral description’ (Guillot 2017, p. 34). Construed as such, it makes very little sense to argue that a phenomenal experience may persist devoid of a subject because the notion of phenomenal consciousness entails subjective character and thereby a subject.

Whilst this argument seems reasonable enough, it is relatively straightforward to spot that there is a false equivalency at work here, for just because we cannot conceive of phenomenal experience without a subject does not entail that we cannot conceive of a subject without phenomenal experience. On this view, one may realistically reiterate that phenomenal experiences are relational and thereby maintain that phenomenality only arises because of the relation between the subject and some content. I suggest that the problem with this argument is twofold, however. Firstly, it seems possible to bring the motivation for the relational account into doubt by constructing an argument to ground the contention that to doubt one’s existence as an experiential subject is itself a manifestation of the experience of doubt, and therefore one cannot reasonably attest that the subject may exist devoid of experience because the act of doubting one’s experiential content is itself affirmation of the content being doubted. Construed as such, one might argue that to deny one’s existence as a subject of experience is seemingly entirely self-refuting and therefore it is impossible, or at least extraordinarily explanatorily difficult, to coherently disentwine the subject from its experiences. I articulate this argument as follows:

1. *If* there is something it is like to doubt one's existence as an experiential subject, then this experience is a manifestation of phenomenal doubt<sup>76</sup>.
2. An experience of phenomenal doubt is the manifestation of a unitarily bounded subjective experience of low informational entropy that meaningfully manifests as a phenomenal experience with a particular subjective character for a given entity S and not all those entities  $\sim$ S. (i.e., phenomenal doubt is an experience with a subjective character that is experientially meaningful to one particular entity).
3. Therefore, an experience of phenomenal doubt is an instantiation of unitarily bounded subjective content that experientially manifests to a particular subject 'S'.
4. Therefore, if there is something it is like to doubt one's existence as an experiential subject, one cannot doubt one's existence as an experiential subject.

In this sense, it is not immediately clear how we might epistemologically justify separating the subject from experiences for it is difficult to justify the denial of the reality of one's existence as a subject of experience in a way that might lead to a positive concept of an 'experience-less-subject'. We might say, therefore, that just as I cannot epistemically separate my experiences from myself, nor can I epistemically separate myself from my experiences.

---

<sup>76</sup> The most obvious way to respond to this premise, and indeed this argument, is in physicalist terms by denying the reality of not just one's initial phenomenal experience but also of both the concept of subjectivity and one's existence as a subject. Such a response would fall neatly in line with Dennett's (1991) illusionism. Whilst I acknowledge this response, I do not devote a great deal of time to this on the grounds that this thesis is not concerned with establishing the reality of phenomenal properties or subjectivity. Instead, as iterated within the introduction of this piece, I am concerned exclusively with motivating panpsychism, which, under any reasonable definition, takes phenomenal properties and subjectivity to be actually existent and non-illusory. I do note, however, that Dennett arguably has never overcome the question: If subjectivity is illusory to such an extent that neither subjective consciousness nor the subject exists, to whom does this illusion present itself?

With this said, there are multiple ways around this argument. One way, as shall be covered shortly, is to hold that whilst I cannot epistemically separate myself from my experiences (or vice versa), perhaps I can ontologically. This argument might hold that an a priori argument could be constructed in which a subject is separated from its experiences, and therefore a positive concept of an experience-less-subject might be achievable (I am personally unconvinced that we can construct such a positive concept, although I note that this does remain an available view)<sup>77</sup>. Another way around this argument may be to reject the initial premise by denying the reality of one's initial phenomenal experience. These responses may be framed in terms of some form of representationalism and/or the phenomenal concept strategy. In these cases, the subject is posited as the purely physical-functional mechanism that grounds awareness, and thereby may seemingly simultaneously exist as a subject – that is, exist as a unitarily bounded instance of subjectivity - even in the absence of experience.

This strategy, however, leads us back to the initial problem with relational accounts (see the previous section) and must confront the problems inherent in those arguments that attempt to provide an ontological argument for the separation of subjects and experiences, for this highlights the difficulty of accounting for precisely how subjectivity occurred if subjective phenomenal character (i.e., subjective conscious experience) is not posited as an ontologically fundamental feature of the universe. In this sense, even if we could somehow conceive of a model for a subject that does not have any subjective content (i.e., a model for a subject that avoids the definition of subjectivity I am using), if the subject may persist devoid of experience, we require an account of precisely how a subject of this kind occurred if it is

---

<sup>77</sup> As intimated, I return to this argument shortly, but suffice to say that the advocate of such a position might arguably still face a substantial difficulty in generating a subject without making use of the notion of subjective conscious experiences (and thereby reinventing the difficulties involved in separating subjects from experiences). See the next section for an account of the difficulties entailed in conceiving of an instantiation of subjectivity devoid of subjective content.

not itself positioned as ontologically fundamental<sup>78</sup>. On a physical-functional account, this would mean explaining precisely how subjectivity ever occurred within a system composed entirely of structural relations. Framed in this manner, this problem is a reiteration of the problem erected against Coleman (2014/17) in the previous section, for just as Coleman faced the issue of deriving subjectivity from that which is subject-less, so too does any advocate of the possibility that a physical-functional subject may emerge from subject-less structural relations. In what follows, I articulate the depths of this problem and note that this issue is no different in kind from the problem facing the idealist cosmopsychism espoused by Kastrup 2018 in section 1.2.3, before concluding that all monistic ontologies that do not coherently incorporate subjectivity at the bottom-level face a not dissimilar explanatory issue.

Consider the world as described by the physicalist. It consists of several structural relations that are themselves wholly, and exclusively, physical (in the sense that they fundamentally instantiate nothing beyond narrowly physical structural properties). Now, consider that the physicalist is wholly constituted of these more fundamental structural relations such that he is decidedly nothing above and beyond a part of this relational universe. Consider now that whilst the physicalist is indisputably a part of the relational universe, he is equally (somehow) able to observe the ostensive reality, and perceivable nature, of the structured universe as an entity that is seemingly distinct from himself. If this is so, as Spencer-Brown (1969) elucidates, ‘we cannot escape the fact that the world we know is construed in order (and thus in such a way to be able) to see itself’ (p. 89), but if this is the case, what is truly astounding

---

<sup>78</sup> I note that there is perhaps a sense that in some dualistic account of reality we might be able to conceive of a subject of this kind by virtue of positing experience-less ontologically fundamental subjects. Although, I do reiterate that I am not clear on how one might do this without incorporating subjective experiences, and I note that in so doing we would also absorb the various explanatory shortfalls entailed by dualism. In this sense, if we could construct a metaphysic that coherently incorporates subjects without inviting the explanatory shortfalls entailed by embracing dualism, this metaphysic would hold an advantage over any dualistic account.



is that the relational universe is capable of perception at all, because in order to achieve this subject-object distinction the universe must have ‘cut itself up into at least one state which sees, and at least one state which is seen’ (p. 89) - in essence, the universe must split itself into a subject and an object. Yet, in so doing, the universe invariably remains itself – that is, it shares a brute identity with itself – and therefore any attempt to legitimise this subject/object distinction as anything beyond an illusion is rendered impossible on the grounds that a monistic universe of this kind does not seem to contain the ontological possibility of a state of affairs in which the universe is distinct from itself to the extent that would be necessary to manifest an instantiation of what we may term ‘subjectivity proper’ – that is, what we may coherently term an entity ‘S’ that is unified and bounded to the extent necessary to demarcate it as the instantiation of a subjective locus with subjective content that is ‘properly’ (in the sense of non-illusorily) isolated from all those entities  $\sim S$ . The core of this argument therefore rests on the notion that proper subjectivity demands a subject/object distinction such that the subject X does not share the same mode of being, or ontological identity, as object Y, yet if monist ontologies akin to physicalism are true, and both X and Y share a brute identity by virtue of being constituted of the same structured relata that are not themselves distinct and isolated loci of subjectivity/phenomenality, then there does not appear to be any coherent means of positioning the identity of X as occupying a unified and bounded locus that is properly distinct and isolated from the identity of Y to the extent that seems necessary in order to coherently incorporate non-illusory subjectivity (or non-illusory subjects) into our worldview. In essence, then, to suggest that certain parts of the structured universe may observe the universe as a distinct object is to suggest that the universe can delude itself into the belief that it is not itself. The question, therefore, becomes: precisely how does the universe split itself to make itself distinct from itself, and precisely how is this undeniably illusory sense of subjectivity a coherent account of subjectivity proper? In this sense, we have

a reiteration of the problem facing Kastrup's Idealist Cosmopsychism (with the only difference being that Kastrup perhaps avoided the first of these questions), and the only solution appears to be found in a redefinition of what is meant by subjectivity such that it is (somehow) redefined in terms that do not include proper subjective content or proper subject/object distinctions.

With the nature of this problem articulated, I maintain that this explanatory hurdle is likely to present itself to many monistic ontologies because a large proportion of the accounts of monism will struggle when tasked with accounting for precisely how their monism may incorporate proper subject/object distinctions. This problem is particularly pressing for monistic metaphysical theories of mind, however, because whilst many theories concede to the illusory nature of subjectivity and subsequently absorb the problems this entails (Dennett 1991/3, Kastrup 2018)<sup>79</sup>, many other theories rest on the notion that there is a non-illusory sense in which we, as perceptual subjects, are proper parts of the universe that are different in kind from the objects we perceive, and as such these theories may reasonably be pressed to account for the sort of subjectivity that is only coherent if the ontology allows for the metaphysical possibility of non-illusory subject/object distinctions.

Arguably, the panpsychist account I am offering avoids this problem entirely, as it entails the possibility that a subject/object distinction exists at the bottom-level because distinct experiential subjects and their structural properties exist at the bottom-level. The problem, however, is now one of reconciling an account of subjects that is consistent with the minimalist ontology that underpins panpsychism.

---

<sup>79</sup> See pages 34-40 of this thesis for an articulation of these problems.

To begin to address this problem, let us begin by picturing a singular bottom-level unit of energy bouncing seemingly aimlessly around the environment in which it finds itself. Now, imagine that this bottom-level unit of energy is itself ontologically fundamental and therefore absolutely minimal, and attach to it the concept of equally minimal and fundamental subjective phenomenal consciousness such that this minimal, absolutely simplistic unit of energy is itself a minimal, absolutely simplistic instantiation of subjectivity that manifests as a unitarily bounded, minimal experience of ‘what it is like’ to be this unit of energy. From here, imagine that the concept of a subject is itself defined in this entirely minimal sense to exclusively denote that which exemplifies a unitarily bounded instance of phenomenal subjectivity. If this picture holds, one might say that if these minimal units of experiential energy are themselves the first instantiations of subjectivity, and we are constituted from similar structured experiential units of energy, then the subjectivity we seem to hold is reducible to the subjectivity that manifests at the bottom-level. Whilst a detailed account of how subjectivity may constitute subjectivity is simply a re-articulation of the combination problem that shall be addressed in chapter 5 of this thesis, the more immediate problem is now to define the precise nature of this bottom-level subjectivity by defining the precise nature of bottom-level, minimal subjects of experience.

### 2.2.2 Minimal Subjects of Experience

At first glance, it appears that providing an account of this nature is simply achieved upon providing an account of experiences themselves, for if the subject of experience is simply an instance of experience, the experiencer simply *is* the experience. This seems to

necessarily be the case in so far as it seems extraordinarily difficult to conceive of how we may coherently construe a bottom-level minimal subject as anything more than a simplistic unitarily bounded subjective instantiation of experience, and this seems to motivate the route Strawson (2011/18), Williford (2015) and Zahavi (2005) take upon articulating the notion of ‘minimal subjects’ (or minimal selfhood). For my purposes, I shall be focused on the Strawsonian account, but in all iterations, the subject *is* the experience (or is at least an intrinsic necessary feature of experience) and thereby exists only in so far as the experience exists, such that once the experience ceases, so too does the existence of the subject (Strawson 2018, p. 7).

Strawson (2011/18) notes that whilst this seems to contravene our understanding of ourselves as subjects, and might appear *prima facie* unintuitive, minimal subjects entail a ‘thin subject of experience [that] exists if and only if experience exists that it is the experience of’ (See Strawson 2011, p. 260), and I maintain that this is justified on the grounds that whilst this account is largely at odds with the currently popular notion that ‘subjects are simply “collections of experiential powers” - things that can exist without ever being the source or site of actual experience’ (Strawson 2011, p. 261) - a minimal subject is all that can reasonably be entailed upon instantiating bottom-level phenomenal properties, because 1) If we hold that subjective experiences require subjectivity in order to be realised, then we posit that phenomenal properties entail a subject, and 2) If this is a given, and phenomenal properties are fundamentally ontological basic, then of the possible concepts of the subject that are available, we should embrace the most basic account that best fits our explanatory needs.

Therefore, we might reasonably hold that a minimal subject refers to a unitarily bounded instance of subjectivity, and from here we may ally ourselves with Strawson by maintaining that this minimal subject ‘is identical with its experience’ *and* this experience only exists in so far as there is a ‘living moment of experience’ that exemplifies the shortest period of time in which a particular experiential type can manifest (ibid, p. 263). Construed as such, Strawson offers an account of the subject of experience that *is* the experience itself and maintains that, as this experience itself manifests as a moment in time, so too does the subject associated with this particular experiential type. Therefore, as Zahavi (2005) elucidates, the subject is ‘not something standing beyond or opposed to the stream of experiences but is rather a feature or function of its [the experiences] givenness. (Zahavi 2005, p. 106)’. In this sense, both Zahavi and Strawson maintain that because all phenomenal content is necessarily beholden to a subject, the minimal subject of experience must exist as the bounded and unified locus in which a subjective experience of a particular type replete with particular phenomenal content manifests, such that the experiential type within a given phenomenal state forms the ‘what it is like’ to have this particular phenomenal state and thereby *be* this particular subject of experience<sup>80</sup>.

The problem with this, however, is that if, as Strawson (2011) attests<sup>81</sup>, phenomenal properties instantiate minimal subjects that are, by definition, construed as unified and bounded instances of subjectivity that are indistinguishable from instances of types of experiential content that can only manifest for a given period of time (p. 263), then as the subject is contained within the experiential type, each experiential type must

---

<sup>80</sup> I note that Strawson is rightly firm on this point.

<sup>81</sup> I note here that Zahavi does not explicitly address the period of time in which the subject may exist, but also note that it seems difficult to conceive of a way in which Zahavi may avoid the same issue facing Strawson.

instantiate a distinct subject of experience in such a way that the phenomenal redness instantiated at time 1 (T1) must manifest as an instance of subjectivity (S1) in which a subject has the phenomenal content of ‘red’, whilst a phenomenal blueness instantiated at time 2 (T2) must manifest as a distinct instance of subjectivity in which a different subject (S2) has the phenomenal content ‘blue’. We are therefore left with a potentially infinite myriad of distinct subjects of experience that continuously pop in and out of existence, and thereby face both the difficulty of establishing precisely how we may explain a property that strongly emerges and demerges into and out of existence, and the difficulty of explaining how distinct experiential types replete with fleetingly existent, and distinctly experientially isolated subjects, should ever achieve the unity to manifest as subjects capable of undergoing multiple experiential types simultaneously. As a result, I offer a rethinking of the Strawsonian notion of minimal subjects and propose melding this concept with the more standard notion of ‘internal’ subjects – that is, subjects that *are* the intrinsic nature of spatial, physical properties which stand as internal pillars of experience and (potentially) persist both as long as both the instantiation of physicality *and* the experiential content does.

I take this merger to be well motivated on the grounds that 1) a central motivator for panpsychism is its ability to avoid the issues associated with emergent subjects, and 2) a concept of minimal subjects in which distinct subjects exist only in so far as experiential types exist renders the combination problem substantially more difficult to solve, by virtue of moving the issue from one of accounting for how a myriad of fundamental micro-subjects of experience constitute a singular macro-level subject of experience to one of accounting for how a myriad of distinct experiential types should constitute a singular experience of numerous experiential types. Whilst this may seem like a simple

rephrasing of the same problem, the former in fact contains the possibility of distinct experiential types manifesting in a *singular* subject of experience in such a way that the same subject may conceivably persist long enough to undergo phenomenal red at T1 *and* phenomenal blue at T2, whilst the latter categorically does not. This difficulty arises exclusively as a result of Strawson's contention that experiential subjects persist only in so far as the experiential type persists, for this leaves open the possibility that a new subject of experience is instantiated by each experiential type, and this creates a difficulty in so far as it seems inconceivable that a constant stream of new, experientially isolated subjects should ever interact to the extent that would be necessary to ground a unified subject capable of achieving a phenomenal state that acts as a locus for a myriad of distinct experiential types simultaneously.

Whilst I am therefore sympathetic to the notion that a subject of experience cannot coherently exist devoid of experiential content, I hold that it is not totally necessary to posit that minimal subjects pop in and out of existence. In line with constitutive micropsychism, we might coherently posit that a very simple unitary structure (or unit of matter) is one and the same as a very simple unit of subjectivity, such that, for example, a singular, irreducible unitary structure (such as a single unit of energy – i.e., a quark) *is* a singular, irreducible subject of experience. In this respect, this extraordinarily minimal and simplistic subjectivity persists if *both* the unity of its essential spatial structure persists<sup>82</sup> *and* it is undergoing experiential content. The distinction here is that Strawson

---

<sup>82</sup> The use of 'unity' here might make more sense when contextualised in terms of the arguments offered in chapters 3 and 4. I note that, in some sense, there is always at least some unity to a given subject's structure because there are unified and bounded instantiations of phenomenality that manifest at the bottom-level that are the intrinsic natures of distinct and irreducible micro-level instantiations of 'structured' physicality that manifest at the bottom-level. I also note here that I take it as a given that it is not inconceivable that an instantiation of 'structured' physicality might acquire various properties, and thereby occupy various spatiotemporal shapes, or arrangements, whilst simultaneously occupying a spatio-temporal continuity that grounds its identity through time and these changing shapes/states/structures (i.e., a tennis ball occupies various spatiotemporal states upon

posits minimal subjects as potentially an infinite stream of novel subjectivity that may continuously pop in and out of existence, while I posit that a singular subject of experience is embedded in a unitary structure for as long as 1) the unity of the structure persists *and* 2) the subject undergoes experiential content<sup>83</sup>. In this latter case, if the structure persists devoid of some mechanism by which the subject may undergo experiential types, the content the subject undergoes is simplified (but as shall be made clear in the next section, if the subject can glean content from an intrinsic self-representation of its own structure it might only be truly annihilated if its fundamental structure is), *and* if a carrier of the (potential) for experiential content exists devoid of an experiential unit of matter in which it may manifest, that specific instantiation of experiential content never exists.

To simplify this, we may say that as soon as there is something it is like to be a unitary structure of energy, that structure is necessarily a subject by virtue of manifesting as a singular instantiation of unitarily bounded subjective phenomenal content in such a way that there is ‘something it is like’ to be this unit of energy as distinct from all others, and we may qualify this elucidation by asserting that this unitarily bounded subjectively experiential unit of energy will persist as long as *both* its fundamental structure persists and it is undergoing

---

being hit but remains the same spherical object at various temporal points). Because a spatio-temporal continuity of this kind is largely uncontroversial (and perhaps even a ‘logically necessary condition of the identity of physical objects’ (Langtry 1972, p. 184)), I do not devote any time to a defence of the spatio-temporal continuity thesis here. However, as mentioned, I do hint at a way in which we might make sense of this (at least when applied to phenomenality) when I come to advance a solution to the combination problem in chapter 4.

<sup>83</sup> Whilst this shall be expanded upon in the subsequent section, I note that a subject can sustain its own experiential content internally by virtue of the self-representation that manifests as a continuous phenomenal field that might bridge the gap between, for example, the redness of T1 and the blueness of T2 (an account of this phenomenal field shall be explicated in the following section). In this sense, I am *not* positing that the subject, or its intrinsic experiences, are relational.



phenomenal content<sup>84</sup>. With this established, I suggest there are three explanatory benefits to this account. Firstly, it resolves the issues that seem to occur when discussing other accounts of subjects by establishing an account of subjectivity that simultaneously conceivably contains its experiential content necessarily but can also conceivably undergo different types of experiential content without demerging into and out of existence<sup>85</sup>. Secondly, as shall be highlighted in the subsequent sections, by framing subjects as unitary structures, we may posit that the extent of a subject of experience's conscious content and causal powers is dependent upon the structures the subject has access to. In this vein, we might say that as bottom-level subjects of experience occupy minimal spatial points, their experiences are equally minimal and simplistic. Further (once the combination problem has been solved) because the brain is a relatively complex structure, we might expect for the subject associated with this brain to be equally complex (in this sense, the disparity between minimal and higher-order subjects might be made transparent). Thirdly, we avoid the combinatory issue facing Strawsonian minimal subjects by virtue of providing a locus of experience that is not contained within specific experiential types but is instead contained within spatially extended instantiations of physicality, and therefore we might suggest that, in certain iterations (and, again, once the combination problem has been resolved), these subjects may conceivably undergo multiple experiential types simultaneously by manifesting as a unitary, complex, unified and bounded experiential, structured entity (i.e., a complex phenomenal structure that has various parts) in which

---

<sup>84</sup> As shall be explicated in the subsequent section, the precise nature of the relationship between a subject's structure and a subject's experiential content is of substantial importance to my wider account of phenomenal properties.

<sup>85</sup> In this sense, the same instantiation of subjectivity persists even if undergoing distinct types of experiential content as long as all of this content continuously makes up 'what it is like' to be the instantiation of physicality that this subjectivity is the intrinsic nature of. An account of how this might be the case is forthcoming in the next section but is expanded upon within chapters 3 and 4.

these types may coalesce within the multi-modal content of a singular subject's 'what it is like' experience<sup>86</sup>.

This is not to say, however, that this account contains the ontological possibility of a subject that may persist devoid of experiential content, and as such the core of this account of subjects of experience now rests upon an adequate account of what precisely constitutes or causes this content. As such, the coherency and existence of phenomenal properties, and thereby my account subjects of experience, now rests upon an account of what precisely the subject is an experiencer *of*, and, subsequently, a more substantive account of precisely where and how experiential content arises (and ultimately a more complete account what experience *is*). I now turn to provide such an account.

### **2.3. “Experiences-of”: On experiential content and the weak transparency of phenomenal properties**

Attempt to imagine a phenomenally conscious subject in a void-like state in which it is entirely devoid of all experiential content, such that there is decidedly *nothing* 'it is like' to exist as this particular instantiation of subjective conscious experience. I take it that upon attempting to form this mental image one faces a considerable difficulty, and I take it that this is because our best concept of phenomenal consciousness denotes any entity “S” in which there is a contentful experience that manifests as a feeling that there is 'something it is like' to *be* this particular entity 'S' as S undergoes the phenomenal content instantiated in this particular state. As such, I suggest that subjective experience is necessarily experience-for a

---

<sup>86</sup> As shall be made clear in the subsequent chapters, this sentiment should be qualified as follows: a singular subject might undergo a multiplicity of experiential content simultaneously *if* it occupies a structure that is sufficiently complex (i.e., has parts that might reasonably be unified into a singular experiential whole of low informational entropy).

subject *and* experience-of qualitative content, because, for an experience to coherently be ‘like’ anything at all, there must be an experience of *something* that this experience is about or directed towards that is itself contentful enough to instantiate the feeling of ‘what it is like’ for a subject to undergo this content. In this sense, I am positing that it is impossible to conceive of an experience entirely devoid of content, and subsequently, I am asserting that for an experience to *be* an experience it must have some phenomenal content that grounds the specific, subjective spike of feeling that enables an experience to be phenomenally ‘like’ anything at all. I take this to encompass a standard account of both experiences and subjects of experience, and maintain that *if* this account is disputed, one must provide an alternative conception of phenomenality that invariably redefines the concept so as to refer to something other than that associated with subjectively undergoing phenomenal content. With the necessity of the relation between phenomenal properties, subjects of experience and experiences of qualitative content established, I am in a position to formulate the fundamental contention that runs through this particular chapter:

1. If phenomenal properties are anything at all, they are experiences.
2. Experiences must have content in order to be experiences at all.
3. Experiences must have experiencers in order to be experiences at all.
4. Therefore, phenomenal properties must be *experiences-of-content* and *experiences-for-subjects* in order to be phenomenal properties at all.

With this articulated, I now turn to offer an account of how phenomenal properties might achieve their experiential content and an account of precisely what the nature of this content might be. To ground this account, I begin with an analysis of experiential content as it occurs at the macro-level of human consciousness in the hope that this might lead us to a theory of

the nature of micro-level experiential content that maintains coherence with our intuitions about subjective conscious experience as it manifests at the macro-level. From here, I outline some potential ways that representational content might be reconciled with phenomenal content, before highlighting a problem that must be overcome in order to advance a coherent account of how bottom-level phenomenal properties achieve representational content. With this established, I move on to offer an analysis of some potential ways *not* to overcome this problem, before finally highlighting a solution that might overcome the shortfalls of the latter iterations in a way that does not necessarily contravene our understanding of how phenomenal content manifests at the macro-level.

### 2.3.1 A problem for representationalist accounts of phenomenal content

At this moment in space-time I am undergoing a myriad of experiential types, but most starkly I experience a chill as a cool breeze sneaks through the jar in the window. In this case, we might identify the feelings of coldness and motion as the content of this particular experiential type, and it is now a custom of common parlance to identify the content of this experience as derivative from the breeze itself. In this respect, we tend to adopt a naïve realism by identifying experiential content as derivative from the actual, mind-independent state of our environment and thereby construe the narrowly physical properties analysed by physics as the content of the experiences of experiential subjects. Whilst such an elucidation provides a coherent account of the origins of experiential content that is undeniably highly intuitive, the problem with this is twofold: firstly, we are epistemologically limited creatures, confined to a state of consciousness that cannot reach outside of itself in order to objectively delineate the extent to which the experiential types we undergo are truly veridical, and secondly to say that experiential content simply *is* the content of narrowly physical properties

is to misconstrue what it means to be experiential, for experiential content seems to be markedly different from the content contained in narrowly physical properties. As a result, there is a tendency to avoid this latter issue, whilst conservatively ‘hedging one’s bets’ about the veridicality of experiential content, by framing experiential content in representational terms, such that this content ‘is the way the subject perceptually represents her environment as being’ (see Crane 2009, p. 456), but is not necessarily an absolutely veridical representation of the environment itself. As a *weak representationalism*<sup>87</sup> of this kind explains the origins of the content of our experiential states without forcing us to embrace any difficult epistemological commitments, this form of representationalism is almost universally adopted, and commits one to the minimal and intuitive inference that phenomenal experiences always have representational content – that is, the content of our phenomenal experiences always represent *something* about some referent, and are thereby directed towards, and at least partly derivative from, a referent of some kind<sup>88</sup>. Construed as such, as the chill sneaks through the jar in the window, it is the way the chill *feels* from the inside that captures ‘what it is like’ to undergo this experiential content, and it is the fact that this chill represents something cold that captures its representational content. The former therefore denotes the phenomenal character of my subjective experiences and captures ‘what it is like’ for me to feel this chill, whilst the latter represents something cold by virtue of my undergoing an experience *of* something with the content capable of instantiating coldness in me.

---

<sup>87</sup> This form of representationalism merely holds that phenomenal states necessarily have representational content, but does not necessarily commit to any ontological claims pertaining to the precise nature of this content (this might be contrasted with pure/strong forms of representationalism which do make ontological claims of this kind).

<sup>88</sup>. I note here that if one rejects weak representationalism, at the very least one must provide a coherent account of experiential content that is non-representational and must do this in such a way that explains precisely how we can distinguish between degrees of veridical and unveridical experiential types. I note that anti-representationalists such as Brewer (2006) may have attempted to articulate an account of how the latter occurs, but in so doing failed to articulate a coherent account of precisely how experiential content initially occurs non-representationally.

Whilst there are various advantages to adopting a representationalist account of phenomenal content, and thereby at least partially construing phenomenal properties as representational properties<sup>89</sup>, the central question in any account of representational content is the extent to which this content is either *wide*, in the sense that it may not exist in isolation and is thereby determined by a relation between properties, or *narrow* in the sense that it is intrinsically fully determined and thereby may persist even in isolation from all other properties. One way to articulate the distinction between these two concepts may be to frame this content in terms of intrinsic and extrinsic properties, such that narrow representational content is an intrinsic property held by a given entity 'X' solely by virtue of the way 'X' is in itself and wide representational content is an extrinsic property held by 'X' by virtue of the way 'X' relates to external entities. In this respect, as Lewis (1983) argues, 'a thing has its intrinsic properties in virtue of the way that thing, and nothing else, is' (p. 197), and contrarily a thing 'X' has extrinsic properties by virtue of, at least in part, the way things external to 'X' are. Framing the distinction between extrinsic/intrinsic properties in this way is therefore broadly consistent with the accounts offered by Lewis (1983), Langton & Lewis (1998), Langton (2006), Francescotti (1999) and various others, and might be more succinctly articulated by suggesting, as Langton & Lewis (1998) do, that an intrinsic property ought to be identified as being 'independent of accompaniment' (p. 334). On this account, if a given intrinsic property is independent of accompaniment it is 'independent of the way the rest of the world is' (Weatherson 2001, p. 365), and therefore might persist even when it is entirely isolated from all other things. With this distinction in mind, Langton (2006) suggests that 'whether a property is relational or non-relational is primarily a conceptual matter: it is relational just in

---

<sup>89</sup> See Chalmers (2004) for an account of the advantages of construing phenomenal properties as representational properties, and for an account of the issues associated with adopting an anti-representationalism. Concisely, the advantage is that it provides a coherent account for the origins of content without resigning us to any unattractive epistemological claims about the nature or origin of this content.

case it can be represented only by a relational concept' (p. 5). In this case, we might say that narrow representational content is an intrinsic property 'because it is compatible with loneliness', and wide representational content is relational [i.e., an extrinsic property] because we need to talk about other things when describing it (ibid, p. 5).

We may, with this definition in mind, thereby posit that mass is an intrinsic property because this property exists even upon being isolated from all other aspects of the universe.

Conversely, we may posit that weight is an extrinsic property because its existence depends on the property of density standing in some relation to something other than itself (in this case the gravitational field). In this sense, we might say that 'X's' narrowly representational content is an intrinsic property of 'X' that stands in no relation to anything other than itself, whilst 'X's' wide representational content is an extrinsic property derived from the relations between X and various things<sup>90</sup> ~X.

Construed as such, when conceiving of how phenomenal properties achieved their experiential content, it seems intuitive to endorse a wide concept of representational content, and thereby attest that our phenomenal states are derived from the relations between ourselves and our environment. This move is substantiated on the grounds that it is difficult to conceive of how we may achieve experiential content without a relation, and as it is commonly held that relations necessitate two or more relata, it is presumed that experiential content is derived from the representational content gleaned from the relation between the properties of a perceiver "S" and the properties of an (extrinsic) object of perception "~S". Adopting this depiction of phenomenal content therefore seemingly forces us to posit

---

<sup>90</sup> 'Thing' here denotes simply the subject of an attribution of properties.

phenomenal properties as, in some sense, fundamentally extrinsic, relational properties, because if minimal phenomenal properties are necessarily experiences-for a subject and experiences-of content to such an extent that we cannot conceive of their existence without satisfying both such conditions, and if 1) we hold that the content of experiences is gleaned by way of representational content and 2) the only way for a property to instantiate representational content is to stand in relation to another (extrinsic) property, then it seems, on the aforementioned definition of intrinsic properties, phenomenal properties simply cannot be wholly intrinsic. If we wish to deny this account (and as I shall argue, we *must* deny this), it appears *prima facie* that we can only reasonably either deny that representational content is exclusively wide (as shall be my aim), hold that extrinsic properties are in some sense ontologically fundamental, or hold that there is an internal contradiction in the concepts of intrinsic experiences-for and extrinsic experiences-of that renders our concept of phenomenal properties unintelligible<sup>91</sup>.

With this said, at first glance, it really does appear to be the case that we simply cannot help but endorse the view of phenomenal properties as extrinsic properties, for it seems absurd to posit that the experiential, representational content that occurs upon the cool breeze sneaking through my window is itself *not* content gleaned by virtue of a relation between myself and my environment, and we are therefore left with construing phenomenal properties as, in some sense, necessarily extrinsic properties. One potential way to reconcile this with bottom-level phenomenal properties (whilst attempting to retain the benefits of an account that identifies phenomenal properties as experiences-for and experiences-of), might be to admit ontologically fundamental extrinsic properties into our worldview. Esfeld (2004) proposes

---

<sup>91</sup> It does not seem that there is *necessarily* any such contradiction, however, and therefore (if possible) we ought to focus on delineating a means of reconciling the explanatory advantages of construing phenomenal properties as experiences-of and experiences-for with a non-contradictory account of experiential content.



that one way to do this is to reject accounts of Humean Supervenience (of the type endorsed by Lewis 1986), in which it is suggested that everything in the universe supervenes on the dispersion of bottom-level intrinsic properties instantiated at various points in space-time (and the interactions between these various fundamental intrinsic properties), in favour of an account that 1) holds to the reality of bottom-level relations and 2) denies that these relations supervene on the intrinsic properties of their relata. In this sense, we might say that there are no intrinsic properties at all at the fundamental level – instead, there are merely fundamental relational properties devoid of intrinsicity. Whilst we might be able to employ this account to explain how bottom level instantiations of phenomenality have their content, this worldview faces two explanatory difficulties (one of a broadly metaphysical nature and the other of a broadly phenomenal nature). The former issue turns on the difficulty of reconciling relations without relata, and this difficulty arises because, as Esfeld (2004) himself notes, it is often taken that 1) relations require relata, and 2) relata must be ‘things in themselves, that is, have intrinsic properties over and above the relations in which they stand’ (p. 2). With this said, even if we were to hold that Esfeld (2004) achieves his aim of overcoming the paradoxical nature of holding that relations do not require relata, it seems that an account of this kind remains difficult to reconcile with my account of phenomenality, on the grounds that whilst a relational account might explain experiential content, it provides no advantage whatsoever when it comes to explaining phenomenal subjects. This is because, as argued previously, it is extraordinarily difficult to conceive of how phenomenal subjectivity might itself be an extrinsic property derived from relations that are themselves not particulars and thereby not intrinsic properties (i.e., various relations that are not themselves intrinsically phenomenal in the sense of being unified and bounded subjects of experience). In this respect, whilst we may be able to hold that relations are sufficient for experiential content, I maintain there would

first need to be a subject in order for this content to exist<sup>92</sup>. Yet, if our ontology is fundamentally relational (i.e., devoid of particulars with intrinsic properties), then it appears difficult to pin down bottom-level subjects on the grounds that a relational view appears inappropriate to pin down the particularity necessary in order to coherently satisfy our concept of a ‘subject’<sup>93</sup>. For the purposes of this thesis, this in itself might be sufficient to reject this particular worldview outright (on the grounds that I am only concerned with motivating a panpsychist worldview). However, for the sake of argument, we might generously extrapolate two means with which the relationist might respond to this problem: 1) if the relationist maintains that there *are*, at the bottom-level, the particulars necessary in order for us to account for subjects, then arguably she violates the central commitment of an exclusively relationalist worldview by virtue of seeming to commit to the reality of concrete relata (and thereby conceivably intrinsic properties), 2) if the relationist denies that particulars exist at the bottom-level but holds that subjects emerge by virtue of subject-less relations, then we require an account of precisely how the subject emerged within this relational account without re-inviting the issues surrounding the magic of strong emergence, and, further, we might even press for this account to explain how the subject is in itself a non-illusory phenomenon (as covered in the previous section). In this respect, the relational worldview appears to lead us back to the initial hard problem of consciousness by virtue of confronting the initial conceivability concerns that occurred amidst attempts to reconcile phenomenal properties with a purely relational worldview and led to the rejection of

---

<sup>92</sup> See the previous section for an account of why this is the case.

<sup>93</sup> Fundamentally, therefore, this is reducible to the issue of accounting for how we can coherently conceptualise subjects without particulars (and as we shall see therefore also the issue of accounting for how an ontology devoid of particulars can give rise to the sort of particularity necessary to satisfy our concept of subjects). See chapter 2 for an expanded account of this latter problem and an outline of why the former issue is so substantive.

ontologies fundamentally devoid of intrinsic natures (i.e., narrow physicalism) in the first place<sup>94</sup>.

A relational account such as Esfeld's (2004) is perhaps not the sort of account we are looking for, then. However, neither is an account in which phenomenal subjects are positioned as the bottom-level intrinsic properties that glean their content by way of a relation between themselves and their environment because this account separates phenomenal subjects from their content in such a way that reintroduces the conceptual problems entailed by conceiving of an instantiation of phenomenality devoid of either a subject or content. This is especially problematic because this account might conceivably incorporate Humean Supervenience to posit that phenomenal properties are, at least in part, extrinsic properties derived from the relations between relata that are intrinsically phenomenal subjects. In this case, even with the conceptual difficulties this entails set aside, this is clearly something that cannot hold, on the grounds that if phenomenal properties are, even in this limited sense, extrinsic properties, they cannot be ontologically fundamental, because ontological fundamentality requires that the fundamental property is basic in a manner that it must be irreducible to anything else. Yet, if phenomenal properties *are* extrinsic in this sense, they are reducible to a relation between at least two or more intrinsic properties, and subsequently this would violate the panpsychist commitment to fundamental phenomenal properties. Further, we achieve no respite by turning to the worldview that intrinsic phenomenal subjects, and the relations between them, all exist at the bottom-level, because if these relations *are not* non-fundamental but *are* extrinsic properties that ground phenomenal content (i.e., remain properties that can only be described by reference to a multiplicity of spatially distinct things in the world), we are still

---

<sup>94</sup> See Seager (2006) for an account of panpsychism motivated by the need to account for the intrinsic natures that cannot be captured by a purely relational account of extrinsic properties.

left with the seemingly intractable paradoxical prospect of an intrinsic instantiation of phenomenality 'X' that can, by virtue of still being an intrinsic property, *conceivably* persist 'independent of accompaniment' without satisfying our concept of what it means to be phenomenal (i.e., without having any phenomenal content)<sup>95</sup>.

As such, I hold that bottom-level phenomenal properties *can* conceivably be relational but cannot *just* be relational in so far as we must be able to describe them without admitting extrinsic relational terms into our discourse. However, they must also be experiences-of-content, and this content does seem to be best accounted for by use of extrinsic properties. So, the issue is one of explaining how fundamental phenomenal properties can achieve experiential content intrinsically, and how this intrinsic content can be reconciled with the intuition that the representational content manifesting in our phenomenal states can be both representative of our environment *and* weakly transparent in the sense that they reveal something about the environment they represent.

In what follows, I highlight the issues involved within the most robust strategies for solving this issue, before advancing a solution that can explain how phenomenal properties may be both intrinsic and representational of their environment. To reach this point, however, it seems prudent to first simplify the core of the problem. I take it, in a not wholly dissimilar

---

<sup>95</sup> One way around this would be to hold that as relations are fundamental, it is not conceivable that anything should ever exist independent of accompaniment, such that the subject is an intrinsic property that has its content fundamentally because extrinsic relations between subjects exist fundamentally. I suggest that this is perhaps one way around the problem, although I note that it is arguably paradoxical to employ intrinsic properties alongside a fundamentally relational worldview (see Teller 1986 for one attempt to do this), and more, if the phenomenal subject is itself intrinsic (but its content is not) in such a way that might lead us to attempt to form a concept of a phenomenal subject as 'independent of accompaniment', we arguably face the same conceptual difficulty of attempting to reconcile a subject devoid of content. In this respect, regardless of whether we adopt a form of metaphysical nonseparability or not, the notion of intrinsic properties seems to turn on the concept of being 'independent of accompaniment' in such a way that if nothing is properly independent of accompaniment, there are no truly intrinsic properties (and if there are no truly intrinsic properties, we return to the difficulty of reconciling an account of purely extrinsic properties with an account of subjectivity).

vein to Chalmers (2004, p. 158), that this problem arises because of simultaneously holding to the following propositions:

1. Phenomenal properties are necessarily contentful.
2. Phenomenal content is representational content.
3. Phenomenal properties are intrinsic.
4. Representational properties are extrinsic.

Construed as such, this seems to be a genuine problem for the concept of phenomenal properties, for we cannot construe phenomenal properties without experiential content that is representative of something, and yet because *prima facie* it is held that an intrinsic property is non-relational, but representational properties, and thereby the contents of experiences, are only gleaned by virtue of a relation, we cannot coherently conceive of this representational content without rendering phenomenal properties something other than intrinsic properties (by virtue of, perhaps, positing them as beholden to the extrinsic, representational properties that arise within the relation between an experiencer 'S' and an object of experience '~S'). The only way out of this inconsistency is therefore to deny one of the four propositions advanced by Chalmers (2004). However, we find no relief in disputing proposition 1, for we cannot construe a phenomenal experience of 'what it is like' to undergo an experience of 'redness', for example, devoid of the red content which may account for what precisely the phenomenal experience is an experience-of. Equally, whilst proposition 2 may be disputed by adopting some form of anti-representationalism, this strategy must involve an as yet unarticulated robustly explicated alternative account of the origins and nature of phenomenal content, and this account must make sense of the seeming reality of the relationship between the contents of our phenomenal states and our environment without

relying on any form of representationalism. Finally, we cannot coherently deny proposition 3 without inviting the unattractive inferences (and conceptual ambiguities) entailed in construing a model of phenomenal properties that posits them as anything other than irreducible and intrinsic. Therefore, it seems all that is left is to deny proposition 4.

With this established, I hold, in line with Chalmers (2004), Hanna and Maiese (2009) and Kriegel (2008-14), that the only way to defuse this dilemma is to posit that *some* but not necessarily *all* representational properties are narrow and thereby intrinsic. Framed in this way, however, the problem with construing phenomenal properties as entailing narrowly representational content manifests immediately: precisely how do phenomenal properties instantiate narrowly representational content without extracting the content of these representations from anything external to themselves? The answer, as I see it, shall be found in the nature of Russellian micropsychism, and shall involve adopting a substantially adapted form of Hanna & Maiese's (2009) and Kriegel's (2011) self-representationalism that leads us to construe basic phenomenal properties as intrinsic narrowly self-representational properties that fundamentally have their content without relation to anything other than themselves. Prior to elucidating the intricacies of this account, however, I first point to some ways *not* to solve this problem by highlighting the issues associated with both Chalmers' (2004) employment of what he terms 'Fregean content' and the accounts advanced by Hanna & Maiese (2009) and Kriegel (2008/11). I begin with Chalmers' account.

### 2.3.2 Three ways *not* to overcome the problem

Chalmers' central motivation for his elucidation of Fregean representational content is an attempt to reconcile our intuitions about our relationship to the environment with the

contention that the representational content associated with phenomenal properties might be narrow. As a result, Chalmers rejects the more standard account of Russellian representationalism, which ultimately falls foul to the issues of wide representationalism by virtue of attributing content to a relation between a subject and properties of objects, and instead favours likening representational content to Fregean ‘modes of presentation’. In this sense, whilst Russellian representationalism renders phenomenal content as necessarily derivative from a wide representation of a property a subject stands in relation to, Fregean content ostensibly avoids this commitment to wide representationalism by virtue of construing a distinction between the properties of a referent, for example ‘a green ball’, and the modes of presentation that give rise to the content of the properties of that particular referent. In this case, Chalmers employs the Fregean distinction between a singular referent, for example ‘the planet Venus’, and the multiple senses entailed by said referent. The argument follows that just as we employ two different concepts, Hesperus and Phosphorus, to denote the same referent (i.e., Venus), we may also extract multiple modes of presentation for the content of this referent. Chalmers continues by highlighting that the totality of the content associated with this referent is itself the totality of the modes of presentation entailed by this referent, such that a ‘mode of presentation is a condition on extension’ (2004, p. 23) in the sense that our concept of Hesperus, for example, is associated with some condition pertaining to what we take the content of Hesperus to be, and this content is itself derived from the various modes of presentation entailed by Hesperus.

Whilst this content seems to be beholden to concepts, Chalmers holds that we can extend this notion to explain the content of experiences. The argument follows that upon undergoing a phenomenal experience of redness, for example, I am attributing the property of red to a particular object, and thereby, according to Chalmers, this experience involves some mode of

presentation that gives rise to the content of the properties of this object. Therefore, upon undergoing this experience of redness, I am representing the property of redness in a certain mode of presentation in such a way that, for it be a red property at all, it must hold some content that I previously associated with the conditions of this colour experience in the past. As such, the mode of presentation, and thereby the content for the property of redness, might be ‘the property that usually causes phenomenally [red] experiences... in normal conditions’ (Chalmers 2004, p. 24). Construed as such, Chalmers notes that ‘one might propose that phenomenal redness is equivalent to the property of having a certain Fregean content (in the appropriate phenomenal way), where this Fregean content involves a mode of presentation such as the property that normally causes experiences of phenomenal redness’ (p. 25).

With this established, Chalmers notes that we may construe a model for Fregean representationalism by attributing the content of phenomenal properties to an internal Fregean content that arises from internal interactions between a subject and the subject’s memories of previous phenomenal experiences of a particular kind. This, at core, seems to typify Fregean content and demarcates it clearly as an epistemic means of mapping relations between phenomenal properties and thereby forming the content of our representations of phenomenal properties. As such, upon representing the property of redness, we invariably employ some Fregean content to represent it as a property that is not (for example, greenness) and thereby represent it as holding some intrinsic, narrowly representational content gleaned by virtue of the similarities in our modes of presentation for the property of red. Chalmers concludes that construed as such, this view ‘is plausibly a sort of narrow representationalism’, and substantiates this because the ‘Fregean contents of a subject’s experiences are determined by the subject’s phenomenal properties, and that the phenomenal properties are intrinsic properties of a subject’ (p. 26).



The problem with this form of representationalism is twofold, however. Firstly, Chalmers, like all advocates of narrow representational, must face the charge that narrow representational content is not ‘real’ content, as Kreigel (2008) elucidates:

*‘Suppose subject S has a concept M with a content C. C is a narrow content only if M’s having C is a non-relational property of S. But C is a real content<sup>96</sup> only if M’s having C is a relational property of S. Therefore, if C is a narrow content, then C is not real content. That is, narrow content is not real content.’ (p. 310).*

Kreigel (2008) notes that there are two ways to respond to this argument. Either reject the claim that real content is only derived from a relation or advance a coherent model for how real content can be derived exclusively from an intrinsic relation (as shall be explicated, I opt for the former). As Chalmers maintains that veridical representations of our environment can be achieved within the confines of Fregean representational content, it seems he is committed to the contention that narrow representational content *is* real. The problem with this, however, is that in holding to the reality of narrowly representational content Chalmers faces the second issue because it seems that Fregean narrow representationalism still necessarily entails some form of a relation between a subject and an environment. Indeed, if Fregean content involves ‘the property that normally causes phenomenally red experiences’ (Chalmers 2004, p. 25), then it seems to have Fregean content at all, we must first rely on a

---

<sup>96</sup> I take the usage of ‘real’ here to be defined as content that may encompass a veridical representation of a given extension. I take it therefore that content cannot coherently be held to be ‘real’ if it is devoid of an extension which may be employed to track its veridicality. As such, narrow content is commonly defined as ‘not real’, because it does not put us into contact with anything that may count as an external extension, and thereby does not put us into contact with the necessary environmental properties that may be employed to assess veridicality. I make this distinction merely to clarify Kreigel’s sentiments in this passage and thereby remain neutral on whether this usage of ‘real’ ought to be endorsed.

sort of wide representationalism for the subject to *initially* stand in the sort of causal relation to such a property that may eventually lead to Fregean content. Chalmers makes no effort to explain how this property may have initially been generated internally, and thereby invariably seems reliant upon at least some form of extrinsic relation to explain how his purportedly intrinsic Fregean representational content occurred in the first place. This is problematic when contextualised in terms of fundamental phenomenal properties because, for the reasons previously outlined, the content of these fundamental properties simply *must* be initially derived entirely without recourse to external relations, and therefore it seems without an account for how Fregean content may be caused *entirely* internally, we cannot coherently hold that the content of phenomenal properties is purely intrinsic, and thereby cannot coherently substantiate Chalmers' appeal to narrow representationalism when employed as a tool to explain the content of bottom-level phenomenal properties. I therefore suggest that any coherent account of narrowly representational content as it manifests in bottom-level phenomenal properties simply must be *entirely* explained intrinsically in such a way that this content may manifest within a property even if this property is entirely isolated. The problem that such an account faces, however, is circumventing the intuition that representational content must initially be derived from a relation that is necessarily predicated upon the existence of at least one entity "S" and at least one other entity "~S". Arguably, there are conceivably only two means of circumventing this intuition: either delineate a model for an *entirely* intrinsic relational form of narrow representational content that does not in any way rely on extrinsic relations to ground phenomenal content, or advance a form of intrinsic, non-relational, narrow representational content that explains away the intuition that 'real' content must necessarily be relational.

Kriegel (2008/11) appears to opt for the former of these two routes by advancing a form of ‘self-representationalism’, which he articulates by appealing to the notion that upon holding up our hand, for example, we represent a hand that is not a distinct entity but is instead an implicit part of what we are. In this sense, as Kriegel argues (2011), the content gleaned from a relation to my arm is not extrinsic because I do not instantiate it in virtue of standing in a ‘relation *to something that does not overlap me*’. (p. 145). With this same sentiment in mind, Kriegel (2011 p.145) offers an account of how subjectivity might achieve its content intrinsically by virtue of framing it in the same way we might frame the relationship between ourselves and our hand, and Kriegel elucidates this account by appealing to the reality of the self-representational content that seems to occur as a relation between proper parts of a whole, in the sense that: ‘M represents itself iff there are states M1 and M2, such (i) M1 is a proper part of M, (ii) M2 is a proper part of M, (iii) M1 bears the right natural relation to M2, and (iv) M2 bears the right representation-transmission relation R to M’ (Kriegel 2011, p. 59). Construed as such, Kriegel's sentiments may be read as not wholly dissimilar to Hanna and Maiese's (2009) articulation of ‘sensorimotor subjectivity’, which employs a not dissimilar relation between parts and wholes to explain how content might be derived from the intrinsic relations between subjective aspects of a whole and those aspects of the whole that account for spatial and temporal structure. Both accounts might therefore conceivably be presented as coherent forms of narrow representationalism that glean content entirely intrinsically by virtue of the relations between the various parts of a whole.

Unfortunately, when appropriately contextualised in terms of the problem at hand (i.e., in terms of the problem of making sense of how fundamental phenomenal properties achieve their content), both solutions fall foul to the same explanatory ambiguity. For, in

both Kriegel's and Hanna & Maiese's cases, we seem to be presented with accounts that separate aspects of a whole into those aspects that are phenomenally subjective intrinsically and fundamentally non-relationally and those aspects that imbue this non-relational subjectivity with content. If this is the case, then it seems the proper parts 'M1' and 'M2' of a unified whole 'M' may be disentwined into those aspects 'M1' that are non-relationally and intrinsically subjective and those aspects 'M2' which ground the relation necessary for subjective content to supervene on these intrinsic, non-relational aspects. Kriegel (2008) appears to substantiate this particular interpretation within his attestation that the narrow content that occurs because of the relation between parts and wholes might be thought of a property of subjects that 'supervenes on their non-relational properties' (Kriegel 2008, p. 306). Subsequently, the problem appears to be that, in the context of bottom-level phenomenal properties, these non-relational properties are rendered ambiguous, because we simply cannot conceive of subjects of experience (experiences-for) without experiences-of content, and therefore it is unclear how we might employ Kriegel's account to coherently construe a truly non-relational model for subjects of experience that may explain how they might manifest as non-relation aspects that are antecedents to the content of their experiences<sup>97</sup>.

Construed as such, we are left with a not unfamiliar issue, for we appear to resign ourselves to an intrinsic relation that presents phenomenal content as content that

---

<sup>97</sup> As Kriegel does appeal to relations between various non-relational properties in order to account for phenomenal content, it must necessarily be the case that these non-relational parts do not contain their content non-relationally, and subsequently it is difficult to conceive of how Kriegel might avoid this issue. Further, Kriegel arguably finds no respite in responding to this charge with the suggestion that these parts simply *are* the wholes (in the sense that, perhaps, phenomenality contains the parts of content and subjectivity) because if this is the case it is not at all clear why we need to appeal to the relations between these various parts in order to account for the content of the whole (i.e., if we cannot ontologically separate phenomenality into the part concerned with subjectivity and the part concerned with content, why do we need a relational account of phenomenal content?).

supervenes on a purportedly ‘contentless’ minimal subjectivity, and thereby we re-invite the question: how can phenomenal subjectivity be construed as a contentless, non-relational intrinsic property if this subjectivity requires (phenomenal) content to coherently manifest as phenomenal subjectivity at all? It seems as much as Kriegel may wish to maintain that our subjective aspects are in some way non-relational, the account offered fails to offer a compelling theory for how we might conceive of the non-relational subjective aspect existing devoid of content, and as a consequence, he arguably fails to offer a compelling model for how phenomenal properties acquire their content by virtue of not addressing precisely how the non-relational phenomenal subjectivity may initially exist to act as the property upon which self-representational properties supervene. Therefore, Kriegel (2008/11) and Hanna & Maise (2009) are left with only two coherent options: either embrace that phenomenal content is self-representational content that supervenes on non-relational phenomenal properties as these non-relational properties stand in various intrinsic relations (but in so doing concede to the contention that if this non-relational property is subjectivity devoid of content we face the seemingly irreconcilable issue of explaining precisely how phenomenal subjects may manifest devoid of phenomenal content), or maintain that phenomenal properties instantiate fundamental subjects of experience that fundamentally have experiential content necessarily intrinsically and entirely non-rationally, but, in so doing, explain how this content is ‘real’.

### 2.3.3 Towards a potential solution

With this latter sentiment in mind, I suggest that such a model might be constructed by reanalysing the commitments underpinning Russellian micropsychism to reveal a

‘structural subjectivity’ that has intrinsically *non-relational*, self-representational phenomenal content that is both real and potentially weakly transparent. I, therefore, hold that as a Russellian micropsychist articulation of phenomenal properties posits them *as* physical properties to such an extent that we may classify that the physical simultaneously just *is* intrinsically phenomenally and extrinsically structural, then we are justified in accounting for structural content in entirely non-relational terms. This is because whilst Kreigel and Hanna & Maise (2009) might conceivably be willing to posit that phenomenal properties may be disentwined from physical properties as to ground the intrinsic relation necessary in order to produce phenomenal content, on the Russellian micropsychist account of phenomenal properties I have advanced in the previous chapters, phenomenal properties are fundamentally *identical* to (broad) physical properties and therefore may *not* coherently be disentwined or reduced to more fundamental, non-relational aspects (or parts).

On this account, then, fundamental physical properties *are* phenomenal properties, and this fundamental physical-phenomenal property is both experiential and structured. Construed as such, it is relatively simple to extrapolate a model for how experiential content is gleaned non-relationally and intrinsically because conceivably the experience itself simply *is* contentful by virtue of occupying a spatial arrangement that is specifically ‘shaped’, or ‘structured’. We might say that as there is a phenomenal property that is identical, for example, to the non-relational, intrinsic property of mass, the phenomenal property intrinsically and non-relationally represents the structure of mass and thereby gleans its experiential content by virtue of being in the structural state of mass. In this sense, fundamental properties have minimal experiential content intrinsically and non-relationally by virtue of instantiating a subjectivity with a particular

spatial structure and this experiential content is thereby a representation of the spatial arrangement of this fundamental property. Kreigel (2014) does not offer such a view, but does note his openness to the possibility that experiential content might occur by virtue of a brute identity of this kind:

*‘It is coherent to maintain that the relevant experiential character [content] is constituted by physical properties after all, due (say) to some brute identity holding between the former and the latter’.* (Kreigel 2014, p. 235)

To make this view clearer, one might think of subjective experience as a pane of glass. In this analogy, the subject is the glass, and the shape, or structure, of the glass, is the content that the subject’s experience has intrinsically. As there is a brute identity between the subject, the experiential content, and the spatially extended structure, we might say that the subject necessarily has its phenomenal content non-rationally because of self-representing its own structural content entirely non-rationally and with absolute immediacy, by virtue of *being* the intrinsic nature of this structure. Construed as such, the subject has its phenomenal content entirely intrinsically and non-rationally precisely because the subject *is* an instantiation of phenomenality that has a shape or structure entirely intrinsically and non-rationally. This content is conceivably absolutely minimal in the sense that it is entirely irreducible to anything more basic, but it remains non-relational, intrinsic representational content by virtue of representing a given spatiotemporal structure that exists even in absolute isolation from all else. It just so happens that the structure it represents simply *is* a given instantiation of phenomenality, and subsequently it is not erroneous to suggest that an instantiation of phenomenality does have non-relational self-representational content by virtue of its

content being a phenomenal self-representation of the spatial arrangement of the structure that it *is*. Therefore, *an instantiation of phenomenality 'X' has content intrinsically and non-relationally by virtue of a direct self-representation of its spatial arrangement, and therefore any change to its structure is a change to its representational content.*

With the analogy of the pane of glass in mind, it is relatively simple to conceive of how this non-relational theory of representational content may be employed to explain the weak transparency of experiential content and thereby explain away Kriegel's concerns that this content is not 'real'. On this account, we might explain strong transparency as the thesis that, upon directing our attention towards experiential content, we see right through it to such an extent that the objects of our perception pass through the glass in such a way that the glass is rendered entirely transparent. Conversely, weak transparency might be framed as the thesis that objects of perception do not pass straight through but are instead reflected upon the glass, such that they reveal the glass whilst also revealing themselves. The former therefore posits that we are only aware of objects of perception and have no means of delineating the intrinsic properties of consciousness, whilst the latter posits that it is difficult, but not impossible, to direct our attention towards these intrinsic properties and thereby delineate the nature of the glass itself. Clearly, it is the latter form of weak transparency that seems to be most relevant to the account of non-relational representational content I am describing, for if the subject of experience has its content intrinsically and non-relationally by virtue of a self-representation of its own structure, then all relational content might be said to supervene upon the intrinsic non-relational content that manifests in minimal structured subjectivity.



In this sense, contrary to Kriegel who posits that relational content supervenes on contentless subjectivity, I suggest that relational content supervenes on intrinsic, non-relational, subjective, spatially structured content in such a way that changes in the structure of a subject may instantiate changes in the type of phenomenal content the subject has. In this sense, phenomenal properties are not directly acquainted with external objects but, by virtue of being structured, potentially stand in a spatiotemporal relation to other such properties that may, when standing in the right relation, instantiate changes in structure that themselves instantiate changes in the structure that phenomenality represents. Therefore, we may say that the interactions between the spatial aspects of distinct fundamental physical-phenomenal properties may give rise to phenomenal changes in the content of a given subject's experiences as a result of inciting changes to the structure this content is a self-representation of, and thereby interactions between structured phenomenality might coherently manifest as *real* representational content that instantiates a subject's 'immediate experience of itself and the universe in one' (Bradley 1893, p. 410). Construed as such, I suggest that this provides an account of representational content that explains fundamental phenomenal content entirely intrinsically and non-rationally, whilst benefiting from remaining consistent with a sort of weak transparency that posits that we are not directly acquainted with externality but may still stand in some relation to structures beyond our own. In this sense, I take this account to be consistent with our intuitions about our relation to mind-independent objects, as espoused in Eddington's (1927) inference that:

*'We are acquainted with an external world because its fibres run into our own consciousness; it is only our own ends of the fibres that we actually know; from those*

*ends, we more or less successfully reconstruct the rest, as a palaeontologist reconstructs an extinct monster from its footprints'. (p. 187)*

Before offering my own explanation of precisely how the fundamentally structured subjectivity instantiated by phenomenal properties reconstructs externality, I now move to articulate a naturalised account of phenomenal properties that leads us to a theory of precisely what phenomenal properties do. Prior to this, however, I deem it prudent to first offer an abridged reiteration of the account of phenomenal properties I have advanced so far and, in so doing, explain the impact this has on our understanding of the physical.

#### **2.4. Plugging the hole in the physical**

If the Russellian-Eddington formulation of Russellian panpsychism holds, then phenomenal properties are simply the intrinsic nature of matter. On this view, as Goff (2019) notes: 'it is consciousness that breathes fire into the equations of physics' (p. 132). Physics simply characterises the nature of the physical 'from the outside' to capture structural properties such as mass or charge, whereas 'from the inside' mass and charge are simple forms of phenomenality. As Goff (2019) articulates, this panpsychism is in no way dualistic, the view is simply that 'those very properties that physical sciences characterizes behaviouristically are, in their intrinsic nature, forms of consciousness...[and in this sense] Eddington's panpsychism does not *add* to our theory of matter, it merely offers a positive proposal as to what matter essentially is' (p. 136).

Construed as such, the appeal of Russellian Micropsychism is apparent, for it seems to find a means of reconciling consciousness with matter in a manner that does not violate our

scientific worldview. The problem with this account, however, has historically been that whilst the motivation to plug the hole in our understanding of the physical with consciousness has been relatively well-received, we have ultimately remained in the dark about what consciousness, and thereby matter, actually is. In this sense, if we find Eddington's proposal convincing, in order to advance a coherent theory of matter, we must advance a coherent theory of bottom-level consciousness; that is, a coherent theory of precisely what precisely phenomenal properties are.

I take it that the arguments on display in this chapter have gone some way to plugging this hole, and I hope that at the close of this chapter we have a more robust understanding of phenomenal properties and subsequently perhaps a more robust understanding of matter. In this respect, I maintain that defining the nature of bottom-level physical properties in terms of structured phenomenal subjectivity as I have done confers two advantages to the panpsychist worldview. Firstly, it offers a coherent account of what the intrinsic nature of matter may be like that does not violate our common-sense understanding of how phenomenal properties might manifest and present. Secondly, by virtue of defining phenomenality as implicitly connected to a given structure, I take it I avoid the charge that the Russellian-Eddington manoeuvre to establish panpsychism ultimately devolves into a manoeuvre to establish idealism. This charge is explicated by Segal & Goldschmidt (2017) as follows:

*'Either narrow physical features are fundamental or they are not. If they are fundamental, then the manoeuvre fails. For now we cannot very easily reject the causal closure of the narrowly physical domain, and systematic overdetermination looms. So, then the manoeuvre does not establish panpsychism. If, on the other hand, narrowly physical properties are broadly physical but not narrowly physical. And if the manoeuvre establishes panpsychism,*

*then the only such features are mental features. So all narrowly physical features reduce to mental features; and all physical features that are not narrowly physical are mental. So all physical features, period, reduce to mental features. And that's idealism, or something near enough.'* (Segal & Goldschmidt 2017, p. 47).

I hold that on my account of phenomenal properties, we remain true to panpsychism precisely because phenomenality cannot coherently manifest devoid of a spatial arrangement, shape, or structure. Therefore, there can be no phenomenality without physicality. In this sense, a narrow conception of the physical is entirely rejected but it does not necessarily hold therefore that 'all narrowly physical features reduce to mental features', for we might say that phenomenality and physicality are dual modes of presentation (one conceived from the outside, the other from the inside) of a singular phenomenal-physical property that manifests as the quality of structured phenomenal subjectivity and realises physical-phenomenal powers as a consequence. I take this to be a middle ground between a narrow and a radically broad conception of physicality (of the sort that leads to idealism) and thereby hold that this entirely avoids the charge that Russellian panpsychism leads to idealism because on this account phenomenal properties are intrinsically structural and are thereby necessarily physical in such a way that it makes no sense to posit that these structures reduce to purely mental, or phenomenal, features.

With this established, I end this chapter by briefly re-stating my account of the nature of phenomenal properties by describing them as necessarily causally significant experiences-for-subjects and experiences-of-content, and re-affirming that this leads us to an account of fundamental phenomenal properties that, by necessity, defines them as instantiations of causally significant, spatially arranged subjects of experience that hold intrinsic, non-

relational, self-representational content upon which weakly transparent, relational, representational content potentially supervenes.

In essence, then, at the bottom level, we have atoms and molecules that exist as instances of structured phenomenality undergoing the steady experiential hum that is instantiated by a non-relational self-representation of their own spatial arrangement, and we might extrapolate that as these structures collide with other such spatial arrangements, this hum is interrupted by various experiential spikes that accompany the change in a given entity's structure. Or, put differently, we have a non-cognitive field of experience that undergoes brief spikes of extraordinarily simple feeling as bottom-level molecules and atoms collide with one another. Subsequently, we have a great wave of background experiences gently humming to their own rhythms of 'what it is like' to be a given spatial arrangement, which, on occasion, change to produce a brief experiential cascade in which the experiential atomic hum intensifies as atoms dance in a shower of experience, like 'fireflies flickering quietly in the night' (Rosenberg 2004, p. 96).

In this respect, whilst an explicit account of precisely 'what it is like' to be a bottom-level entity is perhaps not achievable, the account of phenomenal properties I have offered does provide at least some insight into the nature of the phenomenal properties as they might manifest at the bottom-level, and thereby also provides the briefest of glimpses into how bottom-level experience *might* manifest. To reiterate where we are left at the end of this chapter, then, we might say that the experiences of quarks (or whichever instantiation of physicality turns out to be fundamental) must be:

- 1) Causally relevant 'powerful qualities'

- 2) Instances of unitarily bounded structured subjectivity.
- 3) Instances of phenomenal content that is intrinsically non-relationally realised by way of a direct self-representation of their own physical form.

With this definition as grounding, I take it that I am now able to substantially expand our understanding of phenomenal properties by employing the definition advanced in this chapter to provide both an account of the powers that may be associated with fundamentally structured subjects of experience and an account of precisely how these powers may fit into our scientific worldview. In this sense, if this chapter was concerned with delineating precisely what phenomenal properties are, the subsequent chapters are concerned with finding a place for phenomenal properties in the natural world and delineating precisely what phenomenal properties *do*.

## Chapter 3

### A Place for Phenomenal Properties

With my account of the nature of phenomenal properties in place, the next task is to delineate how phenomenal properties, so defined, may be reconciled with our scientific worldview in a manner that does not in any way violate either 1) the integrity of the account of phenomenal properties on offer or 2) our scientific understanding of the natural world. With this established, I devote the next two chapters to an elucidation of the extent to which positioning causally significant fundamental phenomenal properties at the bottom level may lead us to an account of the time-direction of the universe that neatly finds a place for phenomenal properties without contradicting our naturalised outlook. I take it that the previous chapters have provided sufficient grounding for the account that follows, and I maintain that the aim of this chapter (and the next) is to 1) find a place for phenomenal properties in the natural world, and 2) find a role for phenomenal properties in the natural world that explains precisely what difference integrating these properties into our account of nature might make to our understanding of the universe. To achieve this, the core of the middle portion of this thesis shall span the course of the next two chapters, and shall be split accordingly: chapter 3 shall outline an account of how we might find a place for phenomenal properties in the natural world and will culminate with an account of a ‘general’ role phenomenal properties might play in dispositions observed in physics, and chapter 4 shall build on the conclusions in this chapter to offer an account of a specific role phenomenal properties might play in the natural world. With this in mind, in sections 3.1 and 3.2 I shall explain precisely why a coherent, and explicit, account of the disposition(s) entailed by phenomenal properties is of such importance in the wake of the ambiguities entailed by categorical accounts of phenomenal properties, and I shall employ this to frame the metaphysical landscape by

motivating an account of liberal naturalism predicated upon both a posteriori and a priori considerations. In 3.3, I shall use these inferences to build upon the account of phenomenal-physical powerful qualities I espoused in the previous chapter by offering an ontology of fundamental phenomenal-physical powerful qualities that turns on a bundled theory of phenomenal powers that might be employed to explain how phenomenality might remain causally relevant at the bottom-level. Having found a place for phenomenal properties in the natural world, in chapter 4 I build on the broad account of what powerful phenomenal-physical qualities might do offered in chapter 3 to unambiguously delineate an exclusively phenomenal power, and I employ this to offer a theory of a causal role this power might play in the natural world by advancing an account in which an instantiation of phenomenality might be disposed to phenomenally individuate, and thereby, given the right reciprocal relations, behave like a phenomenal ‘Maxwellian Demon’ (a ‘P-Demon’). Finally, I end the middle section of this thesis by employing this theory to provide an explicit model for how phenomenal individuation might manifest in the natural world, and I achieve this by adopting Bejan’s (1997-2016) ‘Constructal’ account of the universe’s tendency for self-organisation and design that offers a naturalised account of how the potential ‘Demon-like’ behaviour of structured subjects of experience might be reconciled with our scientific understanding.

With this as grounding, by the end of the middle portion of this thesis (i.e., by the end of chapter 4), I hope to be one step closer to providing a complete micropsychist theory of consciousness by unambiguously identifying a place and role for phenomenal properties in the natural world that turns on the adoption of a liberal form of naturalism and offers a theory of powerful phenomenal qualities that culminates in a novel account of the cause of the tendency for certain systems to survive away from thermodynamic equilibrium. Crucially I maintain that this makes no observational (or predictive) difference to the a posteriori model



for reality advanced in the natural sciences, but that this does provide an explicit and naturalizable theory for precisely what difference phenomenality makes to our understanding of the universe that might be employed in the later chapters to provide an account of how phenomenality manifests at the macro-level of consciousness. I take it, therefore, that providing a sufficiently naturalised ontology shall be key to the integrity of such an account, for placing phenomenal properties as a causal ground for this universally observed tendency should not necessarily lead to overdetermination concerns, or violations of any natural laws or scientific observations, but may potentially lead to a different account of the underlying intrinsic causes of the dispositions we observe in the natural world. I begin by framing an account of a liberalised naturalism which may be employed to motivate the account of phenomenal properties I am offering.

### **3.1 Motivating Liberal Naturalism**

It is not uncommon for the Russellian micropsychist account I am endorsing to be justified on the basis that it provides a means of easily naturalising phenomenal properties by positioning them as one and the same as physical properties (in the sense that the physical is intrinsically phenomenal and extrinsically structural) in such a way that they might be admitted into the properties that are the purview of the natural sciences. For the most part, it is difficult to counter the strength of this inference, especially as it provides a neat means of explaining consciousness whilst avoiding the pitfalls of anti-naturalistic metaphysical commitments, which invariably fall-foul to the epistemic and ontological issues that arise upon endorsing the reality of substances or properties that are not grounded in, or reducible to, physical phenomena, and are thereby not obviously reconcilable with the available empirical data<sup>98</sup>.

---

<sup>98</sup> See Strawson (2012) for a survey of the issues associated with anti-naturalism, and an elucidation of the need for what he terms 'real naturalism'.

Indeed, in the previous chapters, I have motivated the Russellian account I am offering by employing the same justifications. The problem with this, however, is that for the most part, the contemporary panpsychist literature has fallen short of providing a sufficiently robust account of what precisely phenomenal properties are, and subsequently, whilst it seems *prima facie* possible that a Russellian micropsychist account of phenomenal properties may hold the key to a naturalised worldview in which non-epiphenomenal consciousness is not excluded, the literature has consistently failed to address the intricacies of precisely how, and in what manner, phenomenal properties may be reconciled with our scientific understanding of the natural world in such a way that their role in the manifestation of the dispositional properties we observe in nature is rendered explicit. In this vein, I take it that the common tendency to describe phenomenal properties as a categorical ground for the dispositions we observe in the natural world is, in principle, relatively well-motivated, but, in practice, overtly vague. This is because naturalising phenomenal properties in this way renders the precise role phenomenal properties play in the natural world largely ambiguous and thereby renders the dispositions of phenomenal properties equally ambiguous. This problem occurs because, as iterated in the previous chapter, upon placing phenomenal properties as the categorical ground of the physical we must confront the issue that physical properties are causally sufficient to explain physical effects. Mass, for example, can be defined either in terms of inertia or gravitation, charge corresponds to the time-invariant generator of a symmetry group, and spin is a form of angular momentum. If such descriptions are sufficient, we are left asking ‘Where does consciousness fit in?’ ‘What work does it do?’<sup>99</sup> Such questions arise because each of the phenomena the panpsychist describes as physical ‘aspects’ of bottom-level phenomenal properties (i.e., mass, charge, spin, et cetera) are themselves sufficiently accounted for by

---

<sup>99</sup> I note that this issue is (informally) explicated by Pigliucci in a dialogue between Goff and Pigliucci (2020), and also note that Frankish (2016) offers a similar critique.

physics, and so it seems by placing phenomenal properties as the categorical ground of physical phenomenon that are themselves sufficiently explained by narrowly physical properties, we are invariably led back to the question: what difference does consciousness make?

The basis for placing phenomenal properties as a categorical ground of the physical is, therefore, suspect, because it is not at all clear that categorical phenomenal properties are not epiphenomenal, and, as a result, a precise account of the dispositional nature of phenomenal properties remains elusive under any model that places them as a categorical ground. In this respect, I hold that to fully unravel the nature of phenomenal properties, we must position phenomenal properties as powerful qualities, because in so doing we may extrapolate their dispositional tendencies in a manner that immediately renders their place in the natural world unambiguous. The most immediate problem this account faces, however, is one of explaining why we should endorse the worldview I am offering over a worldview that results in ambiguous (or even epiphenomenal) phenomenal properties. In this respect, one might reasonably ask: if all such metaphysical frameworks provide identical, naturalised, empirically equivalent outcomes, from where does one motivate the move to endorse one over another?

### 3.1.1 Naturalism and the problem of ontological plurality

*Prima facie* it seems not unreasonable to posit that the problem of justifying the endorsement of one metaphysical theory over another lies at the core of the problems facing naturalised metaphysics and is reducible to the dichotomy that exists between an empirical, naturalistically motivated ontology and a metaphysical framework that, by necessity,

incorporates a priori reasoning to deduce the deep nature of reality. Defined as such, if our aim is to offer the most naturalistically consistent metaphysical framework, the process of maintaining coherence with a posteriori, scientific methodologies and empirical observations must take precedent over the purported theoretical virtues of the metaphysical framework itself, such that regardless of the ontology being advanced, if one's ontology cannot be made consistent with scientific observations, then one's ontology ought to be rejected. In this respect, if two competing theories of naturalised metaphysics are presented, then whichever fits the empirical observations without contradiction ought to be endorsed, and there appears to be no room in our methodology for an analysis of the theoretical explanatory virtues of a given theory. The problem, as is the case in the ambiguous and unambiguous account of phenomenal properties previously discussed, is that if theory A and theory B both fit the empirical observations without contradiction, then metaphysical naturalism, so defined, seems to confer no coherent means of delineating which theory ought to be endorsed.

As Allen (2012) elucidates, the problem with naturalised metaphysics is not necessarily simply reducible to a dichotomy between the *a priori* justifications found in metaphysics and the *a posteriori* justifications that underpin naturalism, for to entirely disentwine metaphysical naturalism from a priori reasoning is to discount what motivated naturalism in the first place. Instead, the problem, as Allen (2012) so defines it, arises as a result of defining metaphysical naturalism as follows:

1. *Metaphysical naturalism considers 'metaphysics to be continuous with, or a branch of science.'*
2. *[Naturalised] metaphysics shares its methodology with science (whatever the methodology of science may be). Naturalized metaphysical theories are only open to*

*revision on logical grounds or upon those which are naturalistically or empirically acceptable.*

3. *Methods of justification which do not meet these standards are excluded from use as methodological tools in metaphysical theorizing.* (Allen 2012, p. 212)

As such, metaphysical naturalism may be defined as an ontological framework that can incorporate a priori methods and justifications if those methods remain continuous with science and its methodologies in a manner that does not lead to any violations of the dispositional properties documented by the natural sciences. However, when defined as such, the problem of delineating between equally naturalised metaphysical theories appears no less transparent, for by predicating naturalised metaphysics on the proviso that it must share the methodology of science, we seem to re-invite the issue of a potential plurality of equally naturalised, and thereby justified, ontological commitments. Allen expresses this latter problem as follows:

1. *'More than one ontological theory fulfils the same explanatory aims.*
2. *If robust realism is true, then one of these theories is the correct, or true one.*
3. *If one takes a naturalistic approach to metaphysics, then there is no method of choosing between theories which: (a) is naturalistically acceptable; (b) does not presuppose the existence of some of the very ontology postulated by the theory.*
4. *If there is no basis for choice, then any decision between theories would be arbitrary, which is not acceptable from the point of view of realism (nor, perhaps, in general)'*  
(Allen 2012, p. 214).

The point to be laboured here is not that the plurality problem occurs by virtue of a

preoccupation with an unachievable pursuit of absolute certainty; instead, the issue is one of a lack of sufficient epistemic tools to delineate the ‘best’ amongst theories that are equally empirically equivalent (and equally continuous with science). In this sense, the metaphysical naturalist might be read as adopting a form of, what Niiniluoto’s (2003) references as, a ‘critical scientific realism’, in which the tenets underpinning scientific realism (i.e., the commitment to a correspondence theory of truth in which ‘truth’ is gleaned by way of the correspondence between concepts and reality) are softened to admit some degree of fallibilism into our account. On this view, our scientific models might be close to the ‘truth’, but we simply do not hold the necessary epistemic tools to ever describe a theory indubitably as anything more than ‘truthlike’ (i.e., a state in which increasingly accurate concepts are advanced to successfully capture more available data but the absolute ‘truth’ is perhaps never achieved). On this reading, Allen’s point is that if our starting point is some fallible kind of realism, then, when framed in terms of delineating the ‘best’ metaphysical theory of reality (and not necessarily delineating which theory is absolutely ‘true’), it is not immediately clear how we might justify the endorsement of one naturalised metaphysic over another.

This is because, as Allen points out, if naturalised metaphysics must share its methodology with science, then, by necessity, one must either posit an ontology that presupposes the existence of that which is yet to be empirically verified, or posit an ontology that *is* naturalised, but in so doing embrace that a competing ontology that is equally continuous with science and its methodologies is no less naturalised and thereby no less justified (i.e., it captures the same amount of scientific ‘data’ by virtue of being constrained to empirical theorising). Naturalising metaphysics in this manner therefore leads to an ontological plurality that arises as a result of a framework that is unequipped with the necessary epistemic tools to delineate which theory is ‘best’ between empirically equivalent, but overtly

ontologically disparate, metaphysical frameworks. In this sense, even if we do admit some degree of fallibilism into our account, it seems we still face a potential difficulty when attempting to choose between naturalised theories of reality.

### 3.1.2 Liberal Naturalism: A solution to the problem

Whilst Allen offers several potential means of avoiding this issue (one of which I shall return to later in this section), for the most part, I maintain that the issue of ontological plurality may be resolved by sharpening (or perhaps ‘broadening’) what precisely is meant by ‘scientific methodology’. Specifically, I propose maintaining some of the core tenets of the scientific methodology whilst admitting more into our account of ‘data’ than can reasonably be admitted within a strictly naturalised scientific account of this methodology (i.e., I suggest that we admit substantive metaphysical theorising into the methodology employed to uncover the world and constrain this theorising such that it must remain empirically but not theoretically equivalent to a scientific account of this ‘data’). As a result, I ground my account of naturalised phenomenal properties by championing a form of Liberal Naturalism that turns on Paul’s (2012) move to position the explanatory strength of competing metaphysical commitments as a means of assessing the robustness of disparate worldviews, and, in line with Paul (2012), I maintain that this account of explanatory strength ought to be employed to delineate between distinct naturalised accounts of our known reality.

Paul frames this account by noting that the methodological continuity of metaphysics and science need not necessarily cause the metaphysician to succumb to the issue of ontological plurality, for in metaphysics, as with science, the theories that underpin our understanding of the world are, at core, models designed to capture the complexity of the available data in a

manner that is both theoretically and explanatorily powerful *and* logically and empirically continuous with the data set as a whole. In this sense, contrary to the standard interpretations, which posit that 1) scientific methodology is exclusively concerned with empirical theorizing and 2) metaphysical methodology is exclusively concerned with a priori theorizing, Paul argues that the core of both metaphysical and scientific methodology is the construction of a model for reality that is explanatorily robust in such a way that the model entails the best possible account of the world. With this established, Paul highlights that both science and metaphysics employ the same methodology: ‘Inference to the best explanation’ (p. 3)<sup>100</sup>. On this account, as Paul notes, both metaphysics and science are methodologically continuous in so far as *both disciplines are concerned with building the most accurate and explanatorily robust theoretical models for our known reality*. The only disparity, then, between metaphysics and science is that metaphysics is concerned with delineating an account of fundamental ontological truths pertaining to reality, and thereby constructs models with this ontology in mind, whereas science takes a far less fundamental approach. With these points in mind, we might extrapolate an account of theory choice and model formation that is methodologically continuous with both metaphysics and science:

*A posteriori model formation:* Theory formation in science, as in metaphysics, involves the construction of empirically grounded and logically consistent models that attempt to capture the best possible explanation for the metaphysical facts observed in the dataset.

*A priori model choice:* If two or more models are empirically equivalent, but ontologically disparate (i.e., if two or more models fit the observable data but hold disparate ontological

---

<sup>100</sup> See Lipton (1991), or Pierce’s (1877) account of abductive reasoning, for other attempts to admit inferences to the best explanation into our scientific methodology.



commitments), theory choice should proceed by way of an assessment of a model's total explanatory power when conceived of in terms of the dataset as a whole. And an assessment of a theory's explanatory power should refer exclusively to the elegance, simplicity, and explanatory scope of the theory in question, such that the value of a given model ought to be derived from the extent to which said model presents certain theoretical and explanatory advantages over its counterparts.

Construed as such, one may infer that a theory's explanatory power is derived exclusively from the extent to which said theory reduces the amount of explanatory impotence residing in the model for the data set. Impotence, so defined, may be employed to denote any model-state in which the theory underpinning this model presents an internal contradiction, fails to account for some sub-set of data, fails to predictively explain some sub-set of data, or, in general, fails to provide the most accurate, informed, precise, and complete description of the data in the context of the data set as a whole. In essence, then, explanatory power is concerned with moving theoretical models from a state of relative informational entropy (i.e., a state in which the possible informational states a model may occupy are numerous and potentially conflicting), to a state in which the model occupies a position of low-informational entropy in which the information contained within the model is rendered less ambiguous and thereby more explicit, more complete, and more explanatorily powerful.

To elucidate this point, let us consider the rolling of four hypothetical 'infinity' dice. Imagine that you have sufficient a posteriori knowledge to make predictions about dice and their behaviour, and so you are shocked to find that upon rolling two of these dice infinitely the dice are themselves unfailingly providing the exact same numerical outcomes (i.e., the same numbers presenting on the same faces consistently). Now imagine swapping these dice,

rolling this new set infinitely and finding that these dice do *not* consistently provide the exact same numerical outcomes. In our theory formation pertaining to the dice, we may extrapolate either a) that the two sets of dice share the same properties and therefore any difference in outcome is explained away as a coincidence, or b) that the two sets of dice do not share the same properties and therefore the differences in outcome are reduced to a difference in an unobserved property of the dice. In either case, according to Paul, we would have satisfied the conditions for a posteriori model formation, for, not unreasonably, both theories posit the reality of the dice and draw logical conclusions based upon the available data. Therefore, we are presented with two empirically equivalent theories with disparate ontological and explanatory commitments, and as a result, according to Paul, we ought to address the explanatory power of both such theories by considering the content of the data set as a whole in an attempt to enter into a cost-benefit analysis of the explanatory benefits of one theory over another (and this analysis must itself incorporate the ontologies on offer into our calculation of total explanatory power). In this case, given our understanding of dice within the context of the dataset, adopting theory ‘b’ confers more explanatory power because positing a difference in properties provides an explanation for the differences in outcome that allows us to make predictions pertaining to this particular dataset that can be employed to expand our understanding of the data set as a whole and explains the behaviour observed in this dataset in a manner that is not internally contradictory or inconsistent with data observed elsewhere. Conversely, theory ‘a’ must appeal to chance to explain away the contradictory outcomes, but in so doing must address and explain the difference in outcome by somehow explaining away the relative unlikelihood of two six-sided dice *infinitely* presenting the exact same numerical outcomes at random when contrasted with the relative likelihood that these same dice would infinitely produce the exact same numerical outcomes if they held a property that disposed them to this behaviour (for example if they were perfectly weighted to

produce exactly similar numerical outcomes). In this sense, whilst we must acknowledge that theory 'a' represents a perfectly naturalizable metaphysical possibility, it would be unreasonable to infer that theory 'a' represents the inference to the best possible explanation given our understanding of the data set as a whole. Therefore, to endorse theory 'a' over theory 'b' would introduce far more explanatory ambiguity into our understanding of the data (i.e., into our understanding of the behaviour of these particular dice) than is reasonably necessary, and as a result theory b is justified on the grounds that its explanatory impotence is substantially less than its counterparts.

I suggest, in line with Paul, that a not dissimilar cost-benefit analysis ought to be employed to choose between empirically equivalent, naturalised worldviews. Therefore, I maintain that explanatory power provides the basis of choice that is sufficient to soften the issue inherent in Allen's (2012) account of ontological plurality, if Allen's initial premise: 'more than one ontological theory fulfils the same explanatory aims' (p. 214) might, from the liberal naturalist perspective, now be read to reveal a contradiction in so far as two ontologically disparate theories cannot be of equal explanatory power (by virtue of the disparate ontological commitments themselves now being integrated into the holistic a priori analysis of the explanatory power of a given worldview).

Construed as such, motivating an account of liberal methodological naturalism<sup>101</sup> seems to be relatively simple if we take it that standard accounts of scientific, or non-liberal, naturalism

---

<sup>101</sup> I note that there are numerous ways to conceive of liberal naturalism. In this case, I am construing it as a counterpart to those naturalised worldviews which disallow anything above and beyond empirical theorizing. I note that, on Paul's account, such rigidly naturalistic worldviews would not necessarily be methodologically continuous with science, but liberally naturalised worldviews would be. I also note that liberal naturalism usually entails an ontological commitment motivated by a desire to remain continuous science whilst noting the perceived limits of a purely structural ontology (see p. 148 for a brief articulation of one account of this kind).

identify scientific methodology as one and the same as empirical theorizing to such an extent that two empirically equivalent ontologies are rendered equally explanatorily coherent. In this case, a commitment to this form of naturalism might entail the adoption of modes of enquiry that are described as relatively explanatorily impotent, for it equips us with tools that provide no coherent means of unambiguously justifying the endorsement of one ontology over another. Contrarily, on the liberal naturalist view, one might suggest that the limitations evident within less liberal accounts of naturalism may be overcome, whilst retaining the core tenets of naturalism, by expanding both what is meant by the term ‘natural’ and expanding the theoretical tools employed when investigating the natural. Therefore, whilst a scientific naturalism restricts the scope of ‘scientific methodology’ such that the term cannot reasonably refer to anything above and beyond empirical theorizing, for the liberal naturalist the methodology of metaphysics is continuous with scientific methodology *if*, as Paul suggests, scientific methodology is liberalised to coherently employ *both* a posteriori model formation and an a priori tool for model choice predicated upon an account of explanatory power. Although Allen neither explicitly accepts nor rejects liberal naturalism, this account seems to be broadly continuous with the solution to ontological plurality described as follows:

*‘Metaphysics, on this view, might be closely allied with science, but not methodologically continuous with it: one might still require that such methodology, and the metaphysics which results from its use, does not openly conflict with science, for example. Perhaps one way to achieve this aim would be to permit substantive a priori reasoning about ontology—that is, to permit metaphysical theorizing and conceptual analysis over and above that which would be employed in empirical theorizing—on condition that the resulting theories do not conflict with science, or restrict which of our current scientific hypotheses could turn out to be true.*

Allen (2012, p. 229)

Construed as such, it seems the point of contention rests in the manner in which one defines scientific methodology. In Paul's estimations, a priori reasoning is itself an aspect of scientific methodology, whereas Allen seems to suggest otherwise. Regardless, both seem to arrive at a not dissimilar solution to the problem of ontological plurality in so far as they both concede that one possible means for the naturalist to overcome this particular problem is to embrace a form of naturalism that simultaneously constrains a priori theorizing to that which is physically possible whilst also embracing the additional a priori epistemic tools necessary to choose between empirically equivalent models of reality.

With this account in mind, we are now in a position to preface my metaphysics by adding a condition to the jigsaw of the panpsychist worldview I am constructing. First, an analysis of a given worldview may reasonably include a priori theorising that might be employed as a tool to delineate its total explanatory power. Second, a given worldview must remain empirically equivalent to the scientific, naturalistic worldview that is has been constructed as a result of a posteriori modes of enquiry. To make this point clear, I highlight that I take empirical equivalence, in line with common parlance in this field (see Laudan & Leplin 1991, Acuna & Dieks 2014), to refer to observational (and predictive) equivalence. In this respect, I hold that for any theory A, and any body of observational evidence E, if theory B presents the same set of observational evidence, and can make the same predictions, at time 1 (T1), then A and B are empirically equivalent. I take it therefore that as liberal naturalism purports to be continuous with science, and as science predominantly proceeds by way of empirical observational evidence, a coherently naturalised worldview must at least be observationally equivalent to the model for reality that currently dominates the natural sciences. By this, I mean that whilst we may employ a priori tools to make inferences about the fundamental

cause of the structures found in the natural world, if this cause leads to a state in which our worldview cannot logically or empirically maintain observational equivalence with the structures observed in the natural sciences, then this worldview is no longer coherently naturalised. I, therefore, take this formation to satisfy Paul's (2012) commitment to an a posteriori ground for model formation, and I supplement this by suggesting that *both* empirical (observational) evidence *and* explanatory power count as epistemological justification for a given model of reality. So construed, we may articulate the core of this condition as follows:

*For any given model of reality, model formation should first be grounded in empirical data and from here proceed by way of an a priori analysis of a given model's total explanatory power<sup>102</sup>.*

With this stipulated, I now move to display precisely how this approach may be employed to motivate an account of phenomenal properties as powerful qualities over an account that places these properties as a categorical ground of the physical.

### **3.2 Motivating the move to reject bottom-level categorical phenomenal properties**

On certain accounts (See Rosenberg 2004), a methodological naturalism of the kind I have endorsed is exercised to ground a (liberal) metaphysical naturalism in which a priori theorising about the nature of reality is employed to find a place for phenomenal properties in

---

<sup>102</sup> I note that this is not entirely indistinct from Goff's (2017) articulation of 'analytic phenomenology' but point out that he supplements explanatory power for 'theoretical virtue' (see p. 271) if theoretical virtue is taken to include an account of simplicity, elegance and parsimony.

the natural world. On this account, it is held that whilst fundamental properties may very well be ‘mutually related in a coherent and natural way by a single set of fundamental laws,’ this does not entail that such properties and laws can all be completely captured in physical terms (Rosenberg 2004, p. 9). As Rosenberg puts it: ‘the Liberal Naturalists recognize the possibility that the specifications of physics and what could subsist in a world wholly portrayed by physics may not circumscribe nature’s limits. That allows the Liberal Naturalist to step comfortably outside the standard physicalist ontology while retaining a naturalist outlook’. (Rosenberg 2004, 9). A liberal naturalism of this kind therefore employs additional a priori reasoning to abstract away from a strictly scientific naturalism that exclusively employs empirical methods, and thereby seems to provide a coherent means of reconciling the panpsychist move to posit phenomenal properties as ontologically fundamental aspects of the physical with our scientific worldview. More specifically, in line with the account of phenomenal powerful qualities I ultimately offer, such a theory is not inconsistent with the potential for a non-structural cause of some of the structural relations we observe in the natural world and motivates this on the grounds that as long as this makes no differences to, and causes no conflict with, the empirical data documented in the physical sciences, it is entirely possible that the cause of these structural dispositions is both naturalizable and *not* wholly physical (i.e., not necessarily *narrowly* physical).

From the perspective of Russellian accounts of panpsychism, one potential problem with commitments to potential non-structural causes of the kind described is that the panpsychist account on offer must still overcome the issue of ontological plurality that arises between those panpsychists who endorse the reality of phenomenal properties as categorical grounds and those who posit phenomenal properties as powerful qualities, for both models present equally naturalised empirical equivalence but seemingly no means of epistemologically

justifying the endorsement of one theory over the other. With my account of methodological liberal naturalism now in place, an appeal of explanatory power can resolve this issue by virtue of highlighting the relative power of an account of phenomenal properties, which turns on them being the intrinsic nature of fundamental powerful qualities, when contrasted with a categorical account of phenomenal properties. In this case, it seems the former confers far more explanatory power in so far as an adequately explicated account of the disposition of phenomenal properties would entail an unambiguous account of what precisely phenomenal properties do in the world that can be contrasted with the ambiguity entailed by accounts that appeal to categorical phenomenal properties.

To substantiate this, it is prudent to reiterate that a chief concern associated with placing phenomenal properties as a categorical ground is that this renders panpsychism explanatorily impotent in so far as consciousness is positioned as the fulcrum of all physical processes but is rendered causally redundant as soon as all physical processes are explicable in purely (narrowly) physical terms. On this account, as Frankish (2016) elucidates, ‘panpsychism offers no distinctive predictions or explanations. It finds a place for consciousness in the physical, but that place is a sort of limbo...It places it at the very heart of every physical entity yet threatens to render it explanatorily idle. For the behaviour of subatomic particles and the systems they constitute promises to be fully explained by physics and the other physical sciences’ (p. 1). As I have covered, this critique is troublesome for the panpsychist because whilst it is clear that phenomenal properties *are* causally relevant at the macro-level (see chapter 1), by positioning them as categorical properties we fail to provide any coherent account whatsoever of precisely what difference phenomenal properties make to the world (or at least what difference they make to our understanding of how the world works). As a result, if explanatory power is employed as the tool to delineate between distinct naturalised



accounts, and if an account that positions phenomenal properties as powerful qualities could be advanced that would unambiguously explain the causal relevance of phenomenal properties in a manner that explains what difference they make to the world (or our understanding of it), then this account would be able to explain some sub-set of the data in a way a categorical account seemingly cannot, and would thereby be explanatorily favoured over any account that places phenomenal properties as a categorical ground. We might therefore simplify the core of the argument advanced within this section by stating that a panpsychist theory that can provide a specific account of the dispositions entailed by phenomenal properties is more powerful (explanatorily) than a theory that cannot, and therefore any account that attempts to disambiguate the precise disposition(s) of phenomenal properties is well-motivated on the grounds that such an account simply explains more than an account that renders the behaviour (or role) of phenomenal properties ambiguous.

Subsequently, my task in what follows is to unambiguously delineate a place and role for phenomenal properties in the natural world. As the account I am motivating describes phenomenal properties as powerful qualities, we might think of the nature of this task in terms of delineating precisely which powers manifest as a result of the ways in which the quality of physical-phenomenality exists in the world. To disambiguate this, because I am endorsing a form of Russellian micropsychism and an account of fundamental powerful-qualities, I take it as a given that the manifestations of the dispositions observed at the bottom-level immediately unambiguously reveal the role bottom-level phenomenal qualities play in the world, and this role remains absolutely continuous with the observations and predications in physics by virtue of the quality of phenomenality, and the various ‘aspects’ it

instantiates, *being* the various dispositions observed in physics<sup>103</sup>. However, in line with Rosenberg's (2004) sentiments and my appeal to liberal naturalism, I hold that fundamental phenomenal-physical properties might be disposed to certain ways of being that are not themselves directly observed in physics, and these unobserved ways of being might very well be employed to offer a new explanation for some of the data manifesting in our scientific models for reality by providing a different causal account of some of the behaviours observed in the natural world.

With this in mind, I suggest that the best (and perhaps only) way to naturalise fundamental phenomenal-physical powerful qualities is to begin with the account of phenomenal properties I have offered in chapters 1/2 and from there extrapolate 1) the most reasonable account of the dispositions that may be entailed by these qualities so described, and 2) the most reasonable elucidation of the behaviour that may be made possible by integrating such properties into our worldview. Doing so will require advancing a naturalised ontology of phenomenal-physical powerful qualities and phenomenal causation that can explain precisely what powerful qualities are, precisely what phenomenal-physical powerful qualities are disposed to do, and precisely how these dispositions might manifest. I now turn to provide such an account.

---

<sup>103</sup> Construed as such, we might immediately argue that powerful qualities are more explanatorily powerful than categorical properties, because the former renders the role of phenomenality unambiguous by virtue of the quality of phenomenality unambiguously *being* the various dispositions observed in physics. Conversely, a categorical account cannot as easily hold that the dispositions observed in physics *are* manifestations of phenomenality.

### 3.3 Powerful Phenomenal Qualities: Towards a liberally naturalized ontology of phenomenality

With my articulation of liberal naturalism and my account of phenomenal properties now in place, let us return to the analogy of the infinity dice offered in the introduction of this piece. The dice, in its initial iteration, contained narrowly physical properties, and so, upon being rolled, realised only narrowly physical phenomena. If the account I have offered thus far holds, we might now change the nature of these dice so that they contain broadly physical properties that are extrinsically spatially extended and intrinsically phenomenal<sup>104</sup>, and so upon re-rolling the dice we might say that they now potentially realise both physical phenomena and phenomenal phenomena. This is, I argue, a substantial step in the right direction because our ontological foundation now contains the potential to account for the phenomenal experience that occurs upon attempting to, yet again, make sense of such lofty concepts as infinity dice.

However, these successes are hampered by the explanatory ambiguities surrounding the nature of this new physical-phenomenal ontological foundation, for this structured phenomenality is now, at least for the sake of argument, framed as *the* fundamental powerful quality, and yet we have not established precisely what it means to say that physical-phenomenality is a fundamental powerful quality, precisely how this fundamental

---

<sup>104</sup> To reiterate, on the account offered in chapter 2, bottom-level phenomenal properties are:

- 1) One and the same as physical properties, in so far as they encapsulate bottom-level physical properties when conceived 'from the inside'.
- 2) Powerful qualities
- 3) Instances of unitarily bounded structured subjectivity.
- 4) Instances of phenomenal content that is intrinsically realized by way of a self-representation of their own physical form.
- 5) Therefore, if the above holds, bottom-level phenomenal properties *are* bottom-level physical properties, and this phenomenal-physical property may be conceived of as a causally significant structured/formed/physical subject of experience, or may more broadly be conceived as simply physical-phenomenality/phenomenal-physicality<sup>104</sup>.

phenomenal-physical quality might realise the dispositions observed in physics, and precisely what difference imbuing the physical with phenomenality might make to our understanding of some of the behaviours observed in the natural world. To address these ambiguities, in the remainder of chapter 3, I shall provide an account of powerful qualities that reconciles an identity theory of powers with Russellian Micropsychism and shall employ this account to reveal how this leads us to 1) a theory of how we might reconcile the commitment to Russellian Micropsychism with the commitment to fundamental powerful qualities, and 2) a broad account of the causal role phenomenal properties play. In the subsequent chapter, I shall expand on this account to provide a more explicit theory of precisely what difference imbuing phenomenality into our worldview might make to our understanding of some of the causal forces at play in the universe. To reach this point, I begin with an outline of the nature of powerful qualities.

### 3.3.1 On the nature of powerful qualities

For the advocate of the powerful qualities thesis, dispositionalism might be thought of as what a given property is disposed to do in certain circumstances and qualitativity might be thought of as a ‘matter of how a thing is like’ (Giannotti 2019, p. 1), or what it is to ‘be’ something. Construed as such, we might simplify this by claiming, in line with Strawson (2008), that qualitativity (synonymous with the categorical) is concerned with what a thing is, what a thing is like, or how it is to ‘be’ something as distinct from being some other thing, and so ‘being is categorical [qualitative] being because that’s what it is to be! That’s what being is!’ (p.278)<sup>105</sup>. Conversely, dispositionality is concerned with ways in which a being of a certain type exists, or, more simply, dispositionality is concerned with what a given quality

---

<sup>105</sup> If ‘being’ is here concerned with a description of how it is for a given thing to exist in a non-dispositional sense.

does. So, qualitativity is concerned with the quality of being and dispositionality is concerned with various modes, or ways, that this being might interact with the world.

Described in this way, the appeal of powerful qualities is evident because the quality (or being) is identical with the power (or way of being). As an example, the quality of sphericity is disposed to roll because of being spherical, and we so might say that rolling is one of the ways being spherical might exist in the world. In this sense, the mode of being (or the power) is itself tied to the being to the extent that the identity of the power is one and the same as the identity of the quality disposing the power, such that dispositionality and qualitativity ‘must be thought as unrealizable limits for different ways of being of that property’ (Martin & Heil 1999, p. 46-7). This is one way of conceiving of the ‘identity’ theory of powerful qualities later described by Heil (2003) as follows:

*‘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,  $Pd$  is  $P$ ’s qualitativity  $Pq$ , and each of these is  $P$ :  $Pd = Pq = P$ . (Heil 2003, p. 111)*

In this respect, ‘the qualitative and dispositional are identical with one another and with the unitary intrinsic property itself’ (Martin 2008, p. 65), and so on this view the dispositional, the qualitative and the property all share an identity that is captured by a full account of what it is to be a given property (i.e., how or what it is to have a quality) and the various ways which this being might be disposed to exist in the world (i.e., what this quality does). When applied to fundamental properties, this view is described in various places as the ‘identity theory of powers’ (see Jaworski 2016, Giionnatti 2019, Kadic 2017, Taylor 2018) – the view

that every fundamental property is dispositional and qualitative – and it is precisely this theory of identity that shall be employed to frame the account of fundamental phenomenal-physical powerful qualities I am offering.

### 3.3.2 Russellian Panpsychism and identity theories of powers

On the Russellian Micropsychist account of phenomenal properties I have offered, physical properties are fundamentally intrinsically phenomenal, and so the fundamental properties are structured phenomenality (physical-phenomenality), which are properties that bridge the divide between the phenomenal and physical in the sense that its physicality ‘Pp’ = its phenomenality ‘Pq’. There are various ways that we might reconcile this commitment with the identity theory of powers; we might adopt a standard interpretation that places phenomenality as the intrinsic nature of fundamental powerful qualities so as to posit that all fundamental powerful qualities are in essence phenomenal powerful qualities; alternatively, we might adopt the notion that phenomenal properties are just one type of fundamental powerful quality that exist amongst a myriad of other fundamental powerful qualities (see Kadic 2017 for an account of this type). In what follows, I highlight the difficulties entailed within both such attempts to reconcile panpsychism with the identity theory of powers before offering an account of phenomenal powers that might avoid the unsavory consequences of the former theories whilst retaining the broad benefits of the identity theory of powers and Russellian Micropsychism. I begin with an analysis of Kadic’s position.

In line with the account of powerful qualities I have offered thus far, Kadic (2017) prefaces his account by rejecting the standard Russellian Panpsychist view that phenomenal properties are the intrinsic nature of matter which act as the categorical basis for the dispositions

observed in physics, and instead opts for endorsing an account of an identity theory of powers in which ‘consciousness is a fundamental and ubiquitous property which is at the same time categorical and dispositional; a quality and a power’ (p. 50). At first glance, this seems to be consistent with my own account. However, as we shall see, it is in fact committed to a worldview that reintroduces the problems that an identity theory of phenomenal powers was designed to avoid. In order to reveal the nature of this problem, let us turn to the manner in which Kadic (2017) posits powerful phenomenal qualities:

*‘In identity theory panpsychism, consciousness is a fundamental and ubiquitous property like every other such property (e.g. spin, mass, electric charge, colour charge). In other words, consciousness is a fundamental property whose existence we need to admit in order to explain how complex subjects come into being, not the fundamental property which grounds all others’* (p. 50).

Articulated as such, the immediate problem with this account is that if consciousness is a fundamental powerful quality that stands *in addition to* other basic powerful qualities, it is not at all clear how a fundamentally phenomenal powerful quality would achieve causal traction on reality. For if mass, spin, charge et cetera are all fundamental powerful qualities that are distinct in kind from the fundamental phenomenal powerful quality, precisely what would the quality of phenomenality be disposed to do? There is simply no room in the standard model for an additional causal power of this type, and so, if phenomenality is a power that stands in addition to the dispositions observed in the standard model, Kadic (2017) invariably faces a difficulty in establishing both precisely how the powerful quality of phenomenality might be naturalised and precisely how this new powerful quality might come to interact with other fundamental powerful qualities. It seems either Kadic (2017) must

concede that his theory leads him to a worldview in which the standard model must be revised or must reintroduce the problem the identity theory of powers was designed to avoid by rendering the precise role phenomenality plays highly ambiguous.

In this respect, a more standard Russellian take on the identity theory of powers, if achievable, would hold a distinct advantage in that it would retain the benefits of identifying phenomenal-physical properties as powerful qualities whilst also avoiding the issue that arises upon suggesting that these phenomenal powerful qualities stand in addition to other basic powerful qualities. One way to do this would be to adopt an account similar in kind to the phenomenal powers view offered by Morch (2018), in which it is held that phenomenal states are, or at least metaphysically ground, certain powers. If we were to take this proposition and offer an account in which phenomenality is positioned as the intrinsic nature of all fundamental powerful qualities, we might be left with ‘the intriguing prospect of a micro-idealist view in which all microphysical dispositions and laws are grounded in the distribution of phenomenal states and the phenomenal powers that they ground’ (Chalmers 2018, p. 16-17). Whilst, for my purposes, and the reasons advanced in chapter 1, I ultimately reject Chalmers’ attempt to construe this view as a ‘particularly pure form of idealism’ (ibid, p. 17), I ultimately do opt for the adoption of something very similar to this view. On this account, I commit to the reality of fundamental powerful qualities but hold that phenomenality is the intrinsic nature of each fundamental quality and thereby, because of the identity theory of powers, also each power. In this sense, we might say that the quality of negative charge, for example, is disposed to move towards positively charged objects, and this quality is intrinsically phenomenal. This would lead us to a worldview in which phenomenality is the intrinsic nature of every fundamental dispositional property documented in physics that would retain consistency with Russellian Micropsychism, and avoid the issues



involved in categorical accounts of phenomenal properties, by positing that phenomenality is the intrinsic nature of all of the bottom-level powerful qualities in such a way that phenomenal states just *are* the various microphysical dispositions documented in the standard model. In the strict ontological sense, then, the bottom-level powerful qualities simply *are* bottom-level phenomenal powerful-qualities because phenomenality is what it is to *be* physical and therefore is also what it is to *be* any one of these bottom-level powers.

With this established, I see two potential subtly distinct ways of describing bottom-level phenomenal powerful qualities. We might say either 1) that there is some form of phenomenality that *is*, for example, the powerful quality of negative charge, or 2) that negative charge is *one of* the ways being a given instantiation of phenomenal-physicality might exist in the world. On one reading of the former account, we might hold that this implies that there is a form of phenomenality that corresponds to each fundamental powerful quality, such that there is a phenomenality of mass, spin, colour charge and so forth, to the extent that every single bottom-level powerful quality might reasonably be said to instantiate a distinct mode of phenomenality that is contained exclusively within this instantiation of this particular powerful quality. To disambiguate this, we might think of bottom-level entities, for example quarks, as holding a number of distinct powerful qualities. They might have spin, charge, mass et cetera, and so on this account each of these powerful qualities have an intrinsic phenomenal nature that is alienated (in the same of phenomenally bounded and isolated) from each of the other fundamental powerful qualities. So, the phenomenality associated with mass might be M1, the phenomenality associated with spin might be S1 et cetera, and each of these might be isolated in their own unified and bounded phenomenal fields that have content by virtue of a direction self-representation of the specific quality that they are. Alternatively, on the second account, we might position the quark itself as the

unified and bounded phenomenal field ‘S’ and hold 1) that it itself is a unified and bounded structured phenomenality, and 2) that having mass, spin, charge et cetera are ways of being ‘S’. On this reading, S might be conceived of as a broad powerful quality that is fundamental, and each of its dispositions are simply ways of being this instantiation of structured phenomenality. What this might mean is that each fundamental entity is a distinct type of structured phenomenality that contains its fundamental qualities essentially. So that being a quark just *is* the qualities of mass, spin, charge et cetera, and if quarks just *are* the bottom-level (and the bottom-level just *is* intrinsically phenomenal) then all of these powerful qualities are not themselves distinct instantiations of disparate phenomenal fields but are unified within a number of fundamental, distinct phenomenal fields (i.e., quarks) that contain a multitude of bottom-level qualities necessarily because these qualities are essential to what it is to *be* this particular bottom-level entity (i.e., what it is to be a quark, or whichever entity turns out to be fundamental). If this is so, then having mass, charge, spin, et cetera<sup>106</sup> *is what it is to be* a fundamental entity, and so if these entities are fundamentally phenomenal then these powerful qualities are also fundamentally phenomenal.

To simplify all of this, we might employ the analogy of the sphericity of a ball. Whilst this ball’s sphericity might be identified as a powerful quality of the ball such that this quality may be disentangled from the ball itself, we might suggest that in fact this ball’s sphericity is *essential* to the ball’s existence and therefore the powerful quality of sphericity is what the ball essentially is to such an extent that if we were to posit the ball as fundamental it would make no sense to suggest that sphericity is a quality that is not also posited as fundamental. Now, if this ball’s nature is intrinsically phenomenal, and if there is a brute identity between

---

<sup>106</sup> Or any other fundamental property that is conceivably at the same time categorical and dispositional.

the extrinsic properties of the ball and the ball's intrinsic nature (i.e., the intrinsic nature and the extrinsic nature of the ball are simply two modes of presentation for the same thing, namely, the ball), then it would make no sense to say that the ball's phenomenality is, in this sense, *not* responsible for the ball's capacity to roll downhill because, on this account, the quality of sphericity is one and the same as the quality of phenomenality when conceived 'from the outside'. This is because, in the Russellian-Eddington account of panpsychism I have been endorsing, the ball's extrinsic structure *is* phenomenal in such a way that the powers entailed by the ball being spherical are in essence powers of phenomenality (in the sense that the quality of being spherical is a structural quality that just *is* phenomenal). Similarly, certain qualities might be essential to what it is to be a bottom-level microphysical entity, and as such if a microphysical entity is fundamentally phenomenal, then the identity of these qualities is also fundamentally phenomenal, and subsequently, if the identity theory of powers is true, the powers associated with these qualities are phenomenal powers.

With the disparities between these two accounts now articulated, I hold that we ought to reject the former account in favour of the latter formation, for the former leads us to a worldview in which phenomenality is more widely spread than it reasonably needs to be and, as a result, it adds an unnecessary explanatory and conceptual burden to the combination problem for panpsychism by way of adding more bottom-level subjects to our worldview than is reasonably required<sup>107</sup>. As such, if we wish to avoid making the combination problem any more intractable than it already is, I suggest that we endorse an account of fundamental phenomenal powerful qualities that positions these qualities, and their powers, as essential to

---

<sup>107</sup> In the sense that as soon as we move the issue away from an account of how bottom-level subjects sum towards an account of how bottom-level powerful qualities sum, we introduce a unified and bounded instantiation of phenomenality that corresponds to each individual quality that manifests at the bottom-level and thereby introduce far more subjects than would be introduced in an account which posits these qualities as simply multiple ways of being bottom-level microphysical entities.

the mode of being bottom-level phenomenal-physical entities, and thereby 1) employ the use of fundamental qualities to describe what it is to essentially *be* a bottom-level phenomenal-physical entity and 2) employ the use of fundamental powers to describe how a being of this kind exists in the world. In this respect, bottom-level phenomenal-physical properties *are* all of the qualities of bottom-level microphysical entities, and as such being a bottom-level physical-phenomenal property in this world involves the number of bottom-level qualities (and their powers) necessary to be the bottom-level entities described in the mathematical models underpinning physics. So, as distinct from a categorical account of phenomenal properties which might lead us to the conclusion that ‘fundamental phenomenal properties play fundamental microphysical roles and underlie fundamental microphysical structure’ (Chalmers 2013, p. 10), I hold that fundamental phenomenal properties *are* fundamental microphysical roles that are disposed to all of the ways of being microphysical captured in the mathematical models of microphysical structure.

What this really amounts to is the suggestion that fundamental phenomenal-physical properties are powerful qualities that are all of the fundamental ways of being physical that are described in fundamental physics, by virtue of existing as the intrinsic nature of bottom-level microphysical entities which have a number of qualities necessarily. While this account may appear *prima facie* reasonable, a theory of this kind precipitates a confrontation with a number of seemingly unsavoury consequences. Chiefly, the upshot of this worldview is that structured phenomenality – as a bottom-level fundamental property – seems to manifest as a mode of being that is a number of powerful qualities, and this seems to be something that might violate the identity theory of powers because, on this view, qualities are identical to their dispositions to the extent that the quality of sphericity is identical to the power to roll. And yet, if my account is correct, fundamental qualities such as mass, spin, charge et cetera

are identical to their dispositions *and* identical to a fundamental instantiation of phenomenality. In this sense, we seem to risk violating the identity theory of powers by holding to an identity relation that is binding across decidedly non-identical powerful qualities to the extent that we risk falling foul to the contention that because mass, spin, charge et cetera are all conceivably parts of what it is to be a fundamental instantiation of phenomenal-physicality, we might erroneously hold that the disposition to move towards positively charged objects is identical to the quality of mass (and not charge). In this sense, we need an account that might shed light on precisely how bottom-level powerful qualities can be bottom-level phenomenal-physical properties whilst retaining independent identities (of the sort that appears to be necessary to adequately individuate between them), and also an account of how bottom-level phenomenal-physical properties might *be* numerous qualities and powers without violating the identity theory of powers. I now turn to provide such an account.

### 3.3.3 Towards a bundled theory of fundamental phenomenal powers

One obvious way around the problem of individuating between distinct powerful phenomenal qualities might be to simply endorse the view that there is a distinct form of phenomenality that corresponds to each fundamental powerful quality, so that there is, for example, a unified and bounded phenomenal field associated with spin and a distinct unified and bounded phenomenal field associated with mass. As intimated, whilst this might avoid the individuation problem, I am reticent to embrace this view because it renders the combination problem substantially more complex and intractable by way of moving it away from the problem of accounting for how fundamental microphysical entities combine to the problem of accounting for how the numerous phenomenally isolated fundamental powerful qualities that

make up these microphysical entities combine<sup>108</sup>. Conversely, however, if do we wish to make the combination problem less intractable by embracing a bundled view of the type I ultimately embrace, we are left both with an individuation problem and with the problem of justifying the move to establish fundamental phenomenal relations between bundles of powerful qualities.

The latter of these two problems is relatively easy to resolve given that Russellian Micropsychism does not necessarily commit us to anything more than the claim that whatever turns out to be fundamental just is intrinsically phenomenal. If what is fundamental are concrete particulars that contain various properties, then the intrinsic nature of these concrete particulars is phenomenal. Conversely, if what is fundamental is the properties that make these particulars what they are, then these properties are intrinsically phenomenal. Or, alternatively, if what is ultimately fundamental are concrete particulars that just are, by virtue of a brute identity, some bundle of fundamental powerful qualities that, when bundled in specific ways, form fundamental microphysical entities that manifest as ‘points’, or broadly spatial structures, in space-time, then as the intrinsic nature of this concrete particular (and therefore this bundle of properties) is fundamental, so too is the singular unified and bounded instantiation of phenomenality that is the intrinsic nature of the sum of the bundle as the bundle occupies a fixed arrangement in space-time of the type associated with microphysical entities<sup>109</sup>.

---

<sup>108</sup> In this sense, my concern is not simply that framing the combination problem in this way adds to the scope of the combinations we need to admit into our account. Instead, I am also concerned with the potential explanatory issues surrounding the move from a combination of particulars to a combination of properties.

<sup>109</sup> By means of reconciliation with my account of phenomenal content offered in chapter 2, we might hold that as a given instantiation of phenomenality is the unified intrinsic nature of a given bundle, it is the way that powerful qualities are bundled that *is* the experiential content of the phenomenality that is the ‘connectedness’ underlying this bundle. Such that the direct self-representation of its own ‘form’ or ‘structure’ just is the direct self-representation of the bundle of powerful qualities that it is, and so as the arrangement of its bundle of powerful qualities changes, so too does its experiential content.

The extent to which one endorses any of these accounts is dependent upon the extent to which one can admit of the possibility of bare particulars (i.e., particulars devoid of properties) or universals (i.e., potentially abstract properties), or whether one finds both views dissatisfying. For my purposes, I do not have the space in this thesis to articulate a defence of, or erect an attack upon, the universalist or the advocate of bare particulars. What I can do, however, is motivate a particular view of fundamentality on the grounds that this view 1) at least appears to make the combination problem for micropsychism slightly less explanatorily complex and 2) at least appears to be consistent with my liberally naturalistic methodology in a manner that is no less intuitive, or worthy of consideration, than the former options. On this view, I hold that particulars are fundamental, and particulars are bundles of essential properties, such that if the particular *is* fundamental, so too are the properties that are essential to the ontology of that particular. In this sense, if we were to imagine that a spherical ball was fundamental, we would also be positing the properties of sphericity as fundamental, and we would do this because the property of sphericity appears to be essential to what it is to be a spherical ball. Similarly, if quark-like particulars are fundamental, we would posit the bundle of properties that are essential to quarks as equally fundamental. Now, if this is the view we endorse, we might extrapolate that there is a brute identity between the particular and the bundle of properties that it *is*, such that the essential properties are identical to, and inseparable from, the particular (and vice versa), to the extent that if this particular has an intrinsic phenomenal nature, this phenomenality will act as a fundamental phenomenal relation between all of the properties that are essential to the mode of being of the particular.

Whilst this might appear reasonable enough, one immediate problem with this view is perhaps the inclusion of a fundamental relation at all, because, as Bradley's regress (1893)

seems to reveal, upon establishing a fundamental relation between bundles of properties, we invite an infinite regress. This is because if the relation is independent of its relata, it would require a potentially infinite number of further relations to relate it to the initial relata, and if the relation is internal, it would arguably not be a relation between qualities at all. In response to this, in line with Olson (1987), I hold that we might avoid this infinite regress by simply positing some form of ‘relation’, if the term is used loosely to reference some sort of ‘connectedness’ (See Olson 1987, p. 61), at the bottom-level so as to suggest that the phenomenal connectedness of fundamental powerful qualities is simply a brute fact about the state of affairs that occur at the bottom-level of the universe, which requires no further explanation or reduction. Indeed, given that this is a relatively obvious, although often a somewhat unanalysed consequence, of standard forms of Russellian micropsychism, which enquire as to the intrinsic nature of fundamental microphysical entities (i.e., concrete, but not necessarily ‘bare’, particulars at the bottom-level), I hold that commitment to a fundamental phenomenal ‘connectedness’ that leads us to establish the various powerful qualities of fundamental microphysical entities as simply ‘aspects’, or parts, of ‘what it is like’ to be a fundamental microphysical particular should not be overly controversial.

If this is the case, we are left with a number of fundamental microphysical subjects that *are*, in some sense, the sum<sup>110</sup> of the fundamental bundles of powerful qualities which are the constituents of the universe as we know it. We, therefore, have several concrete points in space-time that are intrinsically phenomenal and are disposed to all of the microphysical

---

<sup>110</sup> Given the commitments to composition as identity that I outline later in the piece, I here take ‘sum’ to simply mean a whole that is the combination of its parts when these parts are taken together. In this sense, the ‘sum’ of a number of powerful qualities is itself a phenomenal powerful quality that *is* nothing over and above the combination of all of the powerful qualities that make up what it is to be a given fundamental instantiation of microphysicality. I note that this might very well distance my view from Olson’s (1987), who seems to, in places, suggest that sums are perhaps more than their parts. Whilst I do note this disparity in our views, I do not deem this distinction to be substantive or pose any issues.



dispositions underlying the mathematical models for the universe. The problem, as intimated, is that this particular account appears to lead us to a difficulty when it comes to individuating between the various powerful qualities that occur as a result of these fundamental instantiations of microphysical subjects. This is because we typically differentiate between powerful qualities by an analysis of either their qualities or their dispositions, and yet, if certain bundles of powerful qualities just are fundamentally essential to fundamental particulars, then 1) we might worry that this leads to a difficulty in delineating precisely which quality or disposition is being analysed if all of these powerful qualities are only analysable in terms of the totality of what it is to be a given fundamental particular, and 2) if two individuals have exactly the same bundles of powerful qualities, these individuals cannot be individuated by appealing to their powerful qualities. In this sense, we have a problem of individuating precisely which quality within a given bundle is involved in a given disposition (or vice versa) and a problem of individuating between two individuals who seem to hold identical bundles of powerful qualities.

Whilst a solution to the former problem might require some additional qualifications and steps, there is arguably a very simple solution to the latter problem that is immediately available to the view I am endorsing. This is because, on this account, instantiations of structured phenomenality are themselves individual points in space-time, and therefore, as Losonsky (1987) points out, we can extrapolate several properties of these individual instantiations that directly 'involve' individuals, because every individual has 'a unique Spatio-temporal location' (ibid, p. 191) that introduces several properties, like 'object Y being 3 meters from object X, or object X being to the left of object Z' for example, that are grounded in individual points in space-time and can thereby be employed to individuate between individuals with seemingly identical bundles of powerful qualities.

In the case of the former problem, we might also employ the fact that structured phenomenality is instantiated in individual points in space-time to extrapolate subtle disparities, or similarities, between the various ways that structured phenomenality is instantiated at the bottom-level (and the various types of bundles that underlie these instantiations). In this case, we might compare the qualities and dispositions that manifest amongst distinct, individual particulars to extrapolate the powerful qualities that make up these various particulars. We might spot that every particular of a certain kind has angular momentum and extrapolate that there is at least one powerful quality that these individual points in space-time share: spin. Contrarily, we might spot that not all kinds with spin also exhibit the generation of symmetry, and extrapolate that spin is not the same as charge. Over time, it is not so much of a leap to surmise that we might employ this same process to make inferences about any number of the individual fundamental powerful qualities necessary to realise the microphysical dispositions underpinning the standard model and, as such, an analysis of the disparities between individual bundles in space-time might very well lead us to a means of individuating between the individual powerful qualities necessary to 1) make these bundles what they are, and 2) make these bundles individual instantiations that are distinct from, or similar to, their neighbours.

With these individuation problems solved, I now turn to recap what we are left with if we endorse the bundled account of phenomenal powers that I have attempted to advance.

### 3.3.4 Where does this leave us?

On the worldview I have ultimately attempted to advance, we have a myriad of fundamental, individual instantiations of phenomenal physicality ‘x, y, z’ that might reasonably all

instantiate disparate bundles of powerful qualities, which might be individuated by an analysis of the individual powerful qualities necessary to be a specific bundle, or an analysis of the differences between what it is to be this bundle 'x' as distinct from a neighbouring bundle 'y' that is instantiated as an individual point in space-time not shared by x. We might, for ease, refer to these various bundles as various 'forms', or 'aspects', of phenomenal-physicality as phenomenal-physicality is distributed at the bottom-level. So, if we endorse the worldview that quarks are the fundamental basis of the universe, we might say that up-quarks, down-quarks, strange-quarks, and any other fundamental variety of quark, all exhibit slight disparities in the manner in which physical-phenomenality manifests fundamentally, but each of these remain instantiations of the property of phenomenal-physicality, and therefore might rightly be described as aspects of this property in the sense that each such entity is an instantiation of unified and bounded structured phenomenality, but each also possesses a realisation of formed phenomenality that instantiates various distinct qualities and dispositions that differentiates it from other realisations that possess a variation in form<sup>111</sup>. These remain individual instantiations of essentially the same category of thing, but how the fundamental powerful quality of phenomenal-physicality manifests is varied based on, for example, the shape, spin, charge, colour charge, mass, spatial location of the instantiation of

---

<sup>111</sup> I note here that on the account I am offering, phenomenal-physicality is a powerful quality instantiated by various fundamental particulars. Each of these particulars is a distinct aspect of phenomenal-physicality that occupies its own unified and bounded phenomenal field that is experientially isolated, and each new aspect of physical-phenomenality that supervenes on these aspects expands our concept of what it is to be phenomenal-physicality. If, therefore, we could achieve access to every possible world we would perhaps have a relatively robust conceptual map of how this property might manifest, but this would still not equip us with the necessary epistemic tools to attest with confidence that this property is, or is not, a universal in the strict ontological sense (because our epistemic framework only extends so far as to reasonably allow us to make inferences pertaining to how this property might manifest in this actual world). As such, I maintain that all talk pertaining to the existence (or lack thereof) of a truly universal phenomenal-physicality is moot. It is sufficient, therefore, to not commit ourselves to any strict ontological claims pertaining to universals, and instead simply attest that the concept of universals, when employed loosely, is a useful tool to guide the discourse concerning the quantity and extent of the instantiations of particulars that seem to be part of this same category— but, as iterated, I make no claim to the existence (or non-existence) of universals in the strict ontological sense.

formed phenomenality, and I hold that these variances were evident at the bottom-level and thereby require no additional explanation

In line with Russellian micropsychism, we might say that these various ‘aspects’, or ‘forms’, of phenomenal-physicality relate to each other in various ways, and various other properties supervene on these aspects which either instantiating another way of being phenomenal-physicality (i.e., a weakly emergent bundle of powerful qualities that are intrinsically phenomenally unified) or a way of being narrowly physical (i.e., a weakly emergent bundle of powerful qualities that are not intrinsically phenomenally unified<sup>112</sup>). With this established, we might even expect, given the account on offer, ‘that certain causal powers are essentially phenomenal powers, and [certain] causal roles could not be played without consciousness’ (Chalmers 2018, p. 49), and therefore we might surmise that even if *hypothetically* it were possible for these weakly emergent bundles to share extrinsically the same sets of powerful qualities, if one is narrowly physical, and the other is intrinsically phenomenal, they might not be capable of fulfilling the same causal roles.

Consequently, I note that as this chapter draws to a close, we are left with two problems: 1) what causes some properties to supervene on phenomenal-physicality to instantiate phenomenal-physicality and other properties to supervene on this property to only instantiate narrowly physical properties? and 2) precisely which powers are essentially, or exclusively, phenomenal powers? To fully address both of these questions, I suggest we should begin with providing an account of a disposition that might reasonably manifest in certain instantiations of phenomenal-physicality, but would not manifest in any instantiation of

---

<sup>112</sup> An account of how this might occur is forthcoming in the later chapters.

narrow physicality, in the hope that we might employ this phenomenal disposition to both 1) expand the explanatory power of the theory of powerful phenomenal-physical qualities on offer by explicitly detailing what difference phenomenality makes to our understanding of the universe, and 2) perhaps lay the foundation for a potential solution to (some of) the combination problems that might explain why phenomenality is occasionally, but not always, lost as the universe reaches higher orders of existence (see chapter 5). With this in mind, I now turn to chapter 4, in which I expand on the account offered in this section in the hope of providing a much more specific outline of an exclusively phenomenal causal role that turns on a theory of a phenomenal power that conceivably potentially *occurs* in all instantiations of phenomenal-physicality but only physically *manifests* in certain ‘aspects’ of phenomenal-physicality *if* phenomenal-physicality is standing in the right reciprocal relation to its environment.

## Chapter 4

### A Role for Phenomenal Properties

By extracting meaning from the symbols amalgamated upon this page, you are delineating the disparities between the blackness of these letters and the white backdrop, and, in so doing, extracting meaning from these symbols by individuating their differences to the extent necessary to demarcate them as words that denote concepts of particular types. This capacity to individuate between these types of complex informational content is quite indisputably a capacity born of higher-order cognition, and therefore a capacity for complex individuation of this sort is *not* implicitly contained within, or exclusively entailed by, our holding phenomenal properties. Whilst this may be the case, it remains difficult to dispute that regardless of whether we hold these higher-order cognitive faculties or not, if we were to remain phenomenal subjects, it is conceivable that we might still be able to differentiate between the phenomenal content associated with the white backdrop and the phenomenal content associated with the black letters, and we might argue that this is possible because this sort of simplistic ‘phenomenal individuation’ is conceivably grounded in ‘what it is like’ to undergo the ‘feel’ of disparate phenomenal content in such a way that is not necessarily dependent on our cognitive faculties. Whilst this particular claim certainly requires more exposition (forthcoming), for now we might say that whilst higher-order cognition might be required to demarcate these symbols as words pertaining to individual concepts of various types, this sort of cognition is not necessarily required to demarcate the differences between ‘what it is like’ to experience, for example, the phenomenal content associated with black or white, and therefore this reliance on higher-order cognition to explain complex individuation does not entail that phenomenal properties categorically do *not* grant the disposition to individuate between simple phenomenal types. With this in mind, we might even argue that it

seems *prima facie* possible that even the phenomenal properties at the bottom-level, and thereby fundamentally contentful subjects of experience, may achieve simplistic individuation by way of phenomenally differentiating the ‘feel’ of being in one experiential state from the ‘feel’ of another *if* standing in the right reciprocal relation to their environment. In this respect, I am suggesting that, in answer to the question posed at the end of chapter 3, it is phenomenal individuation that acts as the exclusively phenomenal power, for, as I shall elucidate, it is this disposition that is grounded exclusively in the phenomenal ‘feel’ of the disparities between ‘what it is like’ to be in various structural states in such a way that this could not reasonably occur in the absence of phenomenality. In this vein, the core of this particular section aims to provide an account of a disposition in which the phenomenal feel of ‘what it is like’ is itself necessary and sufficient to realise the disposition<sup>113</sup>. With this in mind, I shall address the extent to which we may classify a capacity for minimal, phenomenal individuation as a first-order disposition of phenomenal properties and the narrative herein shall take the form of an account of precisely how we may make sense of structured subjects of experience individuating between types of experiential content. I begin with an account of precisely what individuation entails, before moving on to highlight the similarities between this account of individuation and Dessalas & Zalla’s (2011) work on phenomenal consciousness and labelling and finally concluding by sympathetically expanding upon Byrne’s (2004) account of minimal, first-order knowledge in an attempt to advance a case for positing the metaphysical possibility that *all* forms of phenomenal-physicality *can, when*

---

<sup>113</sup> In order to avoid confusion, this is what I mean by an ‘exclusively phenomenal power’ – i.e., a power that is sufficiently accounted for by only and exclusively employing the phenomenal feel of ‘what it is like’ to be a subject. I note that, in line with my account of powerful phenomenal qualities, there is certainly a sense in which all powers are ‘phenomenal powers’. Whilst this is, in some sense, a given, I note that in this case we are often discussing the powers of instantiations of ‘phenomenal-physicality’, such that these powers are grounded in whatever it is be physical (in such a way that if the physical has an intrinsic phenomenal nature, the powers associated with the physical are in some sense phenomenal powers), whereas in the case of ‘exclusively phenomenal powers’ we are discussing powers that are grounded exclusively in the phenomenal ‘feel’ of what it is like to be an instantiation of physicality.

*standing in the right reciprocal relations to their environment*, minimally and simplistically individuate between the types of experiential content they undergo. With this in place, I turn to offer an account of how this disposition might manifest in the natural world by outlining a theory of, what I term, Phenomenal Maxwellian Demons (P-Demons), before expanding the ontology of causation I am offering by establishing precisely where in the natural world phenomenal individuation and P-Demons might fit in.

#### **4.1 On phenomenal properties and Individuation**

Individuation, as the term has been employed in common philosophical parlance, refers to the concept that a given referent 'S' may be identified as an individual entity that is distinct from all those things  $\sim S$ . In this case, when discussing a subject's capacity to individuate between distinct entities, we are referencing a capacity to delineate the differences between S and  $\sim S$  in such a way that identifies S as an entity of a particular type that is decidedly distinct from all those types  $\sim S$ . Traditionally it has been conceded that harbouring this capacity entails knowing precisely what it means for an entity to be both S and  $\sim S$ , and thereby a prerequisite for differentiation, and subsequently individuation, has traditionally entailed reflectively knowing something about the nature of S and  $\sim S$  that reveals, by necessity, an implicit disparity between these entities.

Contrarily, I propose that all 'forms' or 'aspects' of phenomenal-physicality might, when standing in the right reciprocal relation to their environment, minimally individuate between phenomenal content without necessarily knowing anything other than 'what it is like' to undergo S and  $\sim S$ . In this sense, I am in broad agreement with accounts, such as Dessalas & Zalla's (2011), which posit that 'phenomenal consciousness has an adaptive function which is



to allow discrimination and labelling of perceptual and mental states' (p. 8), however, unlike Dessales & Zalla, who frame phenomenal properties as phenotypes that emerge as higher-level phenomenon amidst the evolutionary process, I suggest that 1) phenomenality is fundamentally instantiated at the bottom-level and 2) a disposition entailed by all 'aspects' of this property, regardless of whether this property is supported by cognitive mechanisms or not, is a capacity to simplistically and minimally individuate between phenomenal states (when standing in the right reciprocal relation to their environment). I am therefore necessarily committed to delineating an exclusively phenomenal power that is not in any way reliant upon high-order processes, and as a result, I immediately reject Dessalas & Zalla's (2011) suggestion that phenomenal consciousness is an evolved emergent property disposed to label mental states and guide 'the evolution of cognitive systems towards increasing discriminatory capacities' (ibid, p. 12). I do so not just because it is extraordinarily difficult to establish precisely how phenomenal consciousness might have evolved if our ontology does not coherently incorporate phenomenality, but also because I am fundamentally concerned with advancing an account of a disposition in which the phenomenal feel of 'what it is like' is itself necessary and sufficient to realise the disposition (and therefore I am interested in delineating a purely phenomenal disposition that is entirely grounded in phenomenality and not grounded, in any way, in the wider system or cognitive tools a given instantiation of phenomenality might have access to).

Clearly, with this in mind, it is difficult to reconcile this commitment with the account offered by Dessalas & Zalla, especially as it appears that the processes involved in 'picking out' or 'labelling' phenomenal content is reliant upon a complex form of neural encoding capable of imbuing the subject with the memories and functional tools necessary to pick out and label content as similar, or distinct, from content previously experienced. In this sense, if

we aim to delineate an exclusively phenomenal power, it is not clear that Dessalas & Zalla's account of an emergent phenomenal labelling mechanism is of much use given that it is not immediately clear that this labelling mechanism is exclusively grounded in phenomenal powers at all. I take it, therefore, that Dessalas & Zalla (2011) might be right in so far as their suggestion that phenomenal consciousness may entail a capacity for discrimination of some kind but note that a complex form of labelling is not an adequate candidate for a purely phenomenal disposition because it seems reasonable to maintain that only a phenomenal consciousness supported by a sufficiently complex cognitive system may achieve the capacity to knowingly 'label' or conceptualise such states. In this sense, whilst we might hold that phenomenality certainly has a role in this 'labelling' process (in the sense that 'labelling' might even be described as a higher-order process of discrimination that supervenes on first-order phenomenal individuation processes), it remains the case that labelling that one is in a particular state simply must entail implementing the use of neural encoding (and therefore cannot reasonably be described as an exclusively phenomenal power).

Traditionally, as is the case in the work of Dessales and Zalla, accounts of individuation have been framed in terms of this sort of labelling, and therefore making sense of attempts to 'discard the possibility that it [phenomenal consciousness] is an evolutionary epiphenomenon' (Dessalas and Zalla 2011, p. 10) because 'phenomenal consciousness produced our ability to discriminate mental states' (ibid, p.12) remains difficult given that if individuation and discrimination require knowingly *labelling* one thing as distinct from another, and if the capacity for labelling of this kind is not achieved by only and exclusively employing phenomenality, then the precise role phenomenal consciousness plays in the

evolution of the discriminatory process remains elusive<sup>114</sup>. With this established, I take it that we may elucidate a means of avoiding this problem by establishing a minimal account of individuation that is potentially achieved by exclusively employing first-order, non-cognitive phenomenal consciousness, and in so doing we might both establish a theory of an exclusively phenomenal power and even make sense of Dessales and Zalla's (2011) final inference that 'phenomenal consciousness is what natural selection could act upon' (p. 12). By means of achieving this, I begin by analysing the distinction between two concepts of knowing 'what it is like' to be in a particular phenomenal state.

#### 4.1.1 Grounding an account of first-order phenomenal knowledge

To render the nature of this distinction transparent, in line with Byrne (2004), I offer an articulation of a first-order conception in which knowing 'what it is like' to be in a particular phenomenal state entails nothing over and above simply *being* in that state, and a higher-order conception in which knowing 'what it is like' to be in a particular state is reliant upon a cognitive system capable of providing the encoding to memory that is necessary for the occurrence of an introspective state in which we recognise, and thereby know, what this particular phenomenal state is 'like' as distinct from others. The disparity between first-order and second-order conceptions of knowing 'what it is like' is therefore reducible to the distinction between the assertion that a sufficient condition for knowing what it is like to be in an experiential state S is simply being in a given phenomenal state S, and conversely the assertion that for one to know what the phenomenal experience 'S' is like, one must know

---

<sup>114</sup> Especially as it seems *prima facie* possible that this labelling process could (perhaps) occur in the absence of phenomenal properties as a given system, over time, learns what it means to be in one state as distinct from another. Although, I note that constructing such an account would require overcoming the symbol grounding problem (see Harnad 1990), and this particular problem is something that an account based on phenomenal individuation arguably avoids by virtue of semantic interpretations conceivably being grounded in phenomenal individuation (see chapters 4/6 for an outline of this account).

what past experiences  $\sim S$  were like. The latter is therefore framed as a ‘higher-order’ account precisely because phenomenal consciousness is *not* on its own a sufficient condition for phenomenal knowledge, whilst the former is a first-order account because phenomenal consciousness is itself sufficient for phenomenal knowledge. Construed as such, we may articulate the distinction between these two concepts as follows:

*Higher-Order Phenomenal Knowledge:* A subject knows what it is like to be in experiential state S if the subject is undergoing S, and this is accompanied by the thought that one is experiencing S as distinct from all those experiences  $cS$ .

*First-Order Phenomenal Knowledge:* A subject knows what it is like to be in experiential state S if a subject is undergoing S.

In this sense, if phenomenal knowledge is framed as knowledge that one is in a particular phenomenal state ‘S’, then the higher-order conception seems to lead us to the conclusion that without knowing that one is in state S, one does not truly know ‘what it is like’ to be in S. By means of elucidation, the higher-order theorist might argue that your knowledge of ‘what it is like’ to read the words on this page is dependent upon your capacity to introspectively demarcate this present experience ‘K’ as dissimilar to past  $\sim K$ -like experiences. This, therefore, presupposes that one had various other not-K-like experiences in one’s past, and one can demarcate that these previous experiences were less similar to this present experience than one’s past K-like experiences. Knowing what it is like is therefore reducible to a higher-order labelling process in which incoming experiential data is matched, or not matched, to encoded representations of this data, and therefore without this labelling

process, the subject does not know that one's experience is in state K and subsequently does not know what the state 'K' is like.

As Byrne (2004) elucidates, however, this labelling process entails, at foundation, a capacity to knowingly individuate between phenomenal content by demarcating the similarities or disparities between distinct phenomenal types, and therefore if a sufficient condition for knowing what it is like to be in a particular phenomenal state "K" is just identifying the differences between this state and those states  $\sim K$ , then this is perhaps something that can be achieved by employing purely first-order materials. To ground this point, Byrne asserts that we might think of experiential properties 'P, Q, R' as those properties associated with experiences of green, turquoise and blue, and argues, in a not dissimilar vein to my elucidation of representationalism espoused in the previous chapter, that as experiential content is in part representative of structural relations between a subject and its environment, 'at least *some* of the salient similarities and differences between experiences are due to what they represent about one's environment: because of their content, experiences of green are saliently more similar to experiences of turquoise than to experiences of blue, and so on' (p. 221). As a result, it seems not unreasonable to suggest that as phenomenal disparities/similarities are in part contained in the representation itself, then the phenomenal subject may delineate the similarities between experiences of turquoise and blue without employing any form of introspection by simply being phenomenally experiential of the similarities between the phenomenal experience (or 'feel') of both states, and subsequently thereby may satisfy the sufficient conditions for phenomenal knowledge by phenomenally demarcating the similarities between the sensation of phenomenal green and the sensation of phenomenal turquoise, or even perhaps by demarcating the differences between the sensation of phenomenal red and the sensation of phenomenal blue. With this established, the problem

with higher-order accounts of the type elucidated above is that ‘one’s experience may be like something (in the sense relevant to phenomenal consciousness), even though one does not know that one’s experience is like something’ (Byrne 2004, p. 221), and therefore knowledge that one is in state K is not necessary to know that ‘what it is like’ to be in state K is distinct from ‘what it is like’ to be in state  $\sim K$ <sup>115</sup>.

In this respect, it does not appear overtly difficult to construe an account in which first-order knowledge is achieved by way of exclusively utilising phenomenal consciousness and, from there, employ this idea to provide an account of phenomenal individuation that occurs as a result of a phenomenal subject’s capacity to simply demarcate the phenomenal feel of being in one state from the phenomenal feel of being in another state. In this respect, if a subject S is currently in a phenomenal state K, and this K state is then overlapped by a G state, it appears, given the account of first-order knowledge advanced by Byrne (2004), that S would be capable of phenomenally individuating between the ‘feel’ of the differences in the representation of the K and the G state, and would thereby exhibit some minimal form of first-order knowledge about the phenomenal feel of the state it is currently in. In this respect, S would not necessarily possess any higher-order knowledge at all about ‘what it is like’ to exist in either of these states, but *it would know how it feels to be in these states, and it would therefore be able to simplistically individuate between the disparities in how these states might feel as one state transitions into the next*. Whilst such an account arguably appears intuitive, an account of how a process of individuation of this kind occurs seems to first require an account of precisely how the subject came to exist in a state in which experiential content overlaps to the extent that would be necessary for discrimination of this kind to be

---

<sup>115</sup> Byrne (2004) argues that this realisation occurs as soon as we separate knowing what one’s experience is like from one’s experience ‘being like something’ (p. 222) and maintains therefore that being in a given state K is sufficient to demarcate the phenomenal disparities/similarities between a K state and a not K state.

possible. With this in mind, I now turn to provide an account of phenomenal overlap that invariably also leads to an account of phenomenal continuity.

#### 4.1.2 Phenomenal Continuity and Phenomenal Overlap

Picture yourself awakening this morning and passing through the various stages of conscious experience necessary to reach the point in space-time in which you read this thesis. Suppose that at each stage you are the same conscious subject undergoing disparate experiential content, such that the phenomenal content associated with stirring from one's slumber belongs to the same conscious subject as the phenomenal content associated with the reading of this document. Suppose also that the transition from one's waking to one's reading this document was accompanied by a myriad of other experiences that all seemed to overlap to experientially chronicle the time that amounted between waking to reading this piece. As you underwent these overlapping experiential types, you remained the same subject, and therefore the content of your experiences stood in a co-conscious relation to each other that formed, what Dainton (2000) describes as, 'co-conscious parts of a single multi-modal experience' (p. 79). In this sense, upon waking, you were undergoing a multi-faceted experience of the feel of the bedding, the sounds of your environment, the heat of the air, and these experiences combined to form 'what it is like' for you to awaken. Subsequently, this unified multi-modal awakening experience subtly transitioned as, for example, the phenomenal type associated with waking is overlapped by the phenomenal type associated with an urge to prepare one's breakfast<sup>116</sup>. In this respect, as Pelczar (2014) elucidates, 'when there is a phenomenal overlap between two experiences, the phenomenal content of one of the experiences is

---

<sup>116</sup> To reiterate, if this is to be made consistent with the model I have offered in the previous chapters, we might say that the experiential content, or types, associated with these states are derived from the non-relational self-representations of the various structured brain states the experiential subject undergoes as it exists in the world. In this sense, the phenomenal types associated with these states are individuated by way of belonging to different brain structures that instil distinct phenomenal content.

partially offset towards the future relative to the phenomenal content of the other experience’ (p. 132).

The subtlety of this overlap provides a key insight into the nature of our streams of consciousness, for it seems without an overlap of this kind we are left with an account of conscious experiences that is disparate, disjointed and causally disunified in such a way that one’s experience of awakening is conceivably immediately annihilated and replaced by the entirely unrelated experience of reading this thesis. I hold that an account of this kind ought to be avoided on the grounds that 1) it is not representative of our conscious experiences<sup>117</sup>, and 2) the coherency of a continuous subject of experience is dependent upon a constant, uninterrupted stream of experiential content bombarding the same subject<sup>118</sup>. Therefore, in this latter case, if there are gaps in a subject’s experiential content, the subject that existed upon awakening ceases to exist and is subsequently replaced by an entirely novel subject reading this thesis. As previously articulated, such an account must be avoided on the grounds that it is very difficult to conceive of how an infinite myriad of novel subjects of experience popping in and out of existence in line with new experiences may ever exist in unison long enough to combine into the sort of macro-level conscious subjects we are, and this thereby renders the combination problem far more difficult than is conceivably necessary. With this established, and as a non-sequential account of conscious experience is decidedly not representative of our sequential waking conscious experiences, I take it that phenomenal overlap is necessary on the grounds that it is this that allows us to occupy streams of consciousness that gently

---

<sup>117</sup> Whilst the strength of this statement may be brought into question by pointing out that it falls foul to an anthropomorphic fallacy, I maintain, as I have done throughout, that whilst our conscious experiences may not provide an absolutely steadfast epistemic foundation from which to guess at the nature of bottom-level subjects, our own experiences remain our only means of forming inferences about consciousness at all and therefore ought to act as a grounding for any coherent account of phenomenal properties.

<sup>118</sup> See chapter 2 for an elucidation of why this is the case.



flow from one experiential type to the next in a causal chain of phenomenal content that constantly overlaps in a manner that produces a phenomenal continuity devoid of phenomenal ‘gaps’<sup>119</sup>.

In terms of the representationalist, structural subjectivity I articulated in the previous chapter, a phenomenal overlap of this kind is relatively easy to articulate. Indeed, by framing representational content in terms of a non-relational self-representation of a subject’s structure, and positing that a given subject’s structure may interact (and therefore be changed) by standing in the right relation to external properties, it leaves open the possibility that 1) a subject’s non-relational self-representational content, in some sense, represents external properties by virtue of phenomenally representing the changes external properties instil upon its own structure, and 2) all structural changes<sup>120</sup> that move fundamental subjects away from their base experiential content are represented by the subject as experiential changes that overlap the subject’s base experiential content. In this sense, a given subject S has a minimal state of structural subjectivity that may conceivably persist in isolation from all other entities ~ S, but when standing in the right relation to the structural properties of entities ~S, S’s

---

<sup>119</sup> I note that certain philosophers may argue that an account of this kind must satisfy the ‘bridge problem’ (see Gustafsson 2011). This problem may be articulated as follows: if phenomenal continuity is true, how do we account for things like dreamless sleep? For now, I simply note that, in a strict sense, the subject may conceivably persist in dreamless sleep by demerging into minimal states of phenomenal continuity whilst maintaining a brute identity with itself. Construed as such, one way to address this problem is simply to hold that, in sleep, phenomenality is reduced to the extent that it no longer occupies the structures necessary to recall, report on or willingly access conscious states, but maintain that phenomenality persists in a minimal sense. We might even say that upon waking the phenomenal subject *does* occupy these states and is therefore able to recall, and make sense of, the states it underwent during sleep. Regardless, I shall address this issue in more detail within chapters 5/6.

<sup>120</sup> ‘Structural changes’ here refer to changes in the physical form that the instantiation of phenomenality represents, in the sense that we are discussing changes to the structure that the subject is in. In this respect, as the bundle of properties a given entity has changes, the arrangement of its position in space-time (i.e., its ‘structure’ or physical form) changes, and this is represented in phenomenal changes to ‘what it is like’ to be this entity. Given the way I have employed the concept of ‘structure’ throughout (see p. 9 for my definition), I therefore hold that it is not inconceivable that even quarks should undergo structural changes (and thereby phenomenal changes) as they emit and absorb gluons and undergo various transformations from higher to lower mass states.

structure may change and these changes manifest as phenomenal, representational content that overlaps on S's minimal phenomenal content. This is because, as articulated in the previous chapter, the structure the subject occupies at time 1 (T1) accounts for its minimal, self-representational phenomenal content, which may be said to be its 'base' phenomenal state (i.e., its structure devoid of external relations). Subsequently, when standing in the right relation to external structures at T2, the subject's phenomenal self-representation of its structural content changes in line with the structural changes that occurred upon standing in this relation. In this model, then, the experiential change from T1 to T2 is not an instance of non-overlapping phenomenal content, because we may construe a sequential flow of subjective conscious experience that begins with a subject S's self-representational phenomenal content at T1 which is subsequently overlapped by the phenomenal content that represents the structural changes that occurred in S at T2. In this case, the flow of conscious experience begins with S in a minimal structural, and thereby phenomenal, state, before S undergoes a structural change that subsequently overlaps S's minimal self-representational phenomenal content with something other than the content that manifests in S's base state (or, if we wish to avoid talk of 'base states', simply the content that S was undergoing at T1). With an account of phenomenal overlap and continuity established, I now turn to elucidate precisely how we may combine this model for phenomenal overlap with Byrne's account of first-order phenomenal knowledge in order to explain how subjects of experience, when standing in the right reciprocal relations, are capable of phenomenal individuation.

#### 4.1.3 Reframing Phenomenal Individuation

As articulated in section 3.3.1, one of the problems with my adaption of Byrne's account of first-order phenomenal knowledge was that this account was devoid of an explanation for

precisely how a fundamental subject may phenomenally discriminate between the subjective ‘feel’ of content Y and content Z, if content Y and Z occur at successive instances in time. For without a means of accounting for how Y and Z occur simultaneously within a subject’s conscious experience in a manner that would allow the subject to exclusively employ first-order phenomenal tools to discriminate between them, we necessarily resign ourselves to the inference that if a fundamental subject experienced Y at T1 and Z at T2, and this subject is devoid of an ability to introspectively consult one’s memory of previous Y or Z experiences, then this subject is entirely unable to discriminate between Y and Z and is thereby entirely unable to achieve first-order phenomenal individuation. As a means of avoiding this conclusion, in the previous section I articulated an account of phenomenal overlap, and in what follows I articulate a means in which an overlap of this kind may conceivably be employed to provide an account of first-order phenomenal knowledge, and subsequently first-order phenomenal individuation, in which fundamental phenomenal consciousness is itself sufficient for simple, minimal individuations between distinct types of phenomenal content.

With this established, the core of this account of first-order, minimal, phenomenal individuation may be articulated as follows:

*If knowing what it is like to be in state X is just being in state X, then upon state K overlapping state X, one knows what this overlap is like and may therefore, precisely at the moment in which the overlap occurs, phenomenally individuate between being in this K state and the X state that occurred before the overlap.*

If this holds, then I take it that it is uncontroversial to assert that individuation, of a first-order, minimal and purely phenomenal kind, has occurred, on the grounds that if the

sufficient condition for individuation is a capacity to differentiate between at least two distinct types of content, then this condition is met as soon as a subject (S) undergoes a ‘what it is like’ experience of X that transitions into the subtly distinct experience of ‘what it is like’ for S to undergo K’s content overlapping the content of X<sup>121</sup>. This is because if at time 1 (T1) the subject (S) is undergoing the phenomenal content that manifests in state X, and if, at T2, the phenomenal content of K overlaps the content of X in such a way that, in the moment in which the overlap occurs, the subject invariably knows ‘what it is like’ to undergo this overlap, then the subject (exclusively in the moment of overlap) may phenomenally identify a change in phenomenal content, and thereby may minimally individuate between being in state X and being in the overlapping state. As such, it is this momentary instance of phenomenal overlap that allows that subject to minimally, and purely phenomenally, individuate between two phenomenal states by being in the phenomenal state that manifests in the moment in which one state overlaps the other. This is therefore not cognitively grounded conscious awareness or reflective introspection, but it is individuation of a basic and minimal kind.

With this established, we may frame this overlap in terms of bottom-level subjects of experience by imagining a singular experiential unit of structured energy colliding with another singular experiential unit of structured energy in such a way that both units of energy undergo a change in their structure that manifests as new phenomenal content, which subsequently overlaps the content they were undergoing prior to the collision. In this case, both structured subjects may conceivably be said to individuate between the new state that

---

<sup>121</sup> As articulated in the previous section, it is helpful to think of both X and K as phenomenal content that represents certain distinct, spatially extended structures. In this case, for ease, we might think of ‘X’ as representative of a minimal structure that manifests even upon isolation from all other structures, and ‘K’ as representative of an external, relative structure that supervenes on the minimal structure ‘X’ represents.

occurred upon collision and the previous pre-collision state, and this individuation may be said to occur entirely phenomenally by virtue of both subject's knowing 'what it is like' to be in the overlapping state. We may even streamline this further by suggesting that whenever a structured subject undergoes any interaction with its environment that is sufficient to change or impact its own structure, then the subject will know what it is like for its representation of its new structured state to supervene upon its representation of its base structured state, and will therefore be capable of minimal, first-order phenomenal individuation between the pre-collision representation and the post-collision representation. With this established, I now turn to offer an account of how this behaviour may be naturalised that begins with a theory of one way in which phenomenal individuation may potentially physically manifest, if the dispositional makeup of the universe does not preclude its manifestation.


## **4.2 Towards an account of Phenomenal Maxwellian Demons**

In this moment you are sorting the potentially disparate bits of information contained in your immediate environment into a seamlessly integrated whole in which you, as the subject of experience, are capable of extracting the informational disparities contained within the multi-modal experience you are currently undergoing to such an extent that you are able to demarcate this page, and the content it contains, as different in kind to the various other experiential types that make up your current experiential state. In this process, the human brain may be framed as a sorting machine capable of rapid computations which enable a meaningful interplay between the subject and its environment. Framed in macroscopic terms, this is Maxwell's (1867) 'Sorting Demon' working at such a degree of complexity that it is difficult to reconcile this now evolutionarily optimised system with the sort of sorting demon Maxwell initially had in mind. Yet, when construed correctly, I argue that we may still

employ Maxwell's Demon to reveal the nature of the humble beginnings that ground the complex conscious creatures we are today, and, further, we arguably may achieve this in a manner that emboldens the panpsychist narrative by thinking of Maxwell's sorting demon in phenomenal terms. In what follows, I move to provide an account of Phenomenal, Naturalised Maxwellian Demons (P-Demons) predicated upon the account of phenomenal properties I espoused in chapter 2/3 and the account of phenomenal individuation that arose as a consequence (section 4.1). For the most part, this account will turn on the contention that the disposition entailed by certain 'aspects' of phenomenal-physicality standing in the right relation to their environment is a capacity for phenomenal individuation, and I hold that the behaviour of certain instantiations of structured phenomenality may very well employ phenomenal individuation of this kind to sustain themselves away from thermodynamic equilibrium (and subsequently behave like a crude, naturalised 'Phenomenal Maxwellian Demon'). With this in mind, I begin this account with an outline of the origins of Maxwell's Demon.

Maxwell's Demon began as a thought experiment espoused to reveal a means by which an intelligent agent may potentially violate the second law of thermodynamics by employing information to extract 'free' energy from a system. Maxwell (1867) conceived of an intelligent, micro-level 'Demon' that could track the motion and momentum of every sub-atomic particle in a system and then employ the information of a particles' whereabouts to sort the cold particles from the hot particles by tracking their relative speed. Conceivably, if this were possible, the Demon's behaviour would have violated the second law of thermodynamics by way of decreasing entropy without performing any significant

thermodynamic work at all. We might frame the behaviour of this Demon pictorially as follows<sup>122</sup>:



This thought experiment posed a significant problem for physicists because the second law of thermodynamics states that the entropy of the universe must always increase, and so if Maxwell's hypothetical Demon were possible, the second law could not hold. To grasp the gravity of this thought experiment, and the nature of Maxwellian Demons, it is prudent to elucidate what is at stake when discussing a potential violation of the second law. To achieve this, picture a mug of hot coffee. The second law of thermodynamics captures something implicit about the behaviour of the particles within the coffee: it states that, over time, the heat which occurs because of lots of tightly compacted particles banging together will eventually diminish as the particles contained within this mug dissipate into the wider environment, so that, eventually, the mug and its contents reach thermodynamic equilibrium with the environment in which it is situated. This is the core of the second law of thermodynamics: the universe's entropy (the possible states particles may exist in) always increases, such that the heat within the mug will always diminish over time<sup>123</sup>. Framed as

---

<sup>122</sup> I note that the following was taken from

[https://commons.wikimedia.org/wiki/File:Maxwell%27s\\_demon.svg](https://commons.wikimedia.org/wiki/File:Maxwell%27s_demon.svg) and no changes to the image were made.

<sup>123</sup> Whilst this is one particularly broad elucidation of the second law, I note that this law has historically been articulated in a myriad of ways. Chiefly, when discussing the second law, we are discussing a universal tendency for the universe to move towards states in which 'useful' energy is (in total) lost, as order gives way to disorder.

such, the problem that Maxwell's Demon presents ought to be clear. For, conceivably, if such a Demon were possible, it would be capable of employing information to separate the cold from the hot particles indefinitely to sustain the mug of hot coffee away from thermodynamic equilibrium indefinitely – it would, conceivably, be capable of keeping my mug of hot coffee hot forever!

One way to save the credibility of the second law from this problem was espoused by Landauer (1961), who claimed that whilst it may seem like the Demon is avoiding the inevitable decline into maximal entropy without the exertion of any (significant) thermodynamic work, to function coherently the Demon must store the positions of each individual particle within its memory<sup>124</sup>. If this is the case, as Bennett (1982) showed, the law of maximal entropy remains consistent because the Demon's memory will itself be subject to entropy. What this means is that as the information pertaining to the whereabouts of the particles must be physically encoded somewhere, the Demon itself produces more entropy in the performance of its sorting than it avoids by separating the particles. The crucial move here is to position the Demon as its own source of entropy production which may conceivably allow it to maintain coherence with the second law whilst also slowing the inevitable decline into a state of maximal entropy for as long as its memory lasts.

This version of the Maxwellian Demon might, therefore, be framed as an early iteration of what we now term dissipative disequilibrium systems - that is, systems capable of sustaining themselves away from thermodynamic equilibrium by exchanging energy with their environment in such a way that they propel themselves towards more energetic states and

---

<sup>124</sup> I note that Szilard (1929) proposed a not dissimilar solution upon contending that a real-life Demon would need to have some means of measuring the particles, and the act of measurement would itself require energy (and thereby produce entropy).



thereby, at least for a while, avoid the decline into a maximally entropic state. Construed as such, we are dissipative disequilibrium systems, as are trees, plants, animals, and any other system that sustainably exchanges energy with its environment to maintain its state away from equilibrium. But, crucially, unlike the initial sorting demon espoused by Maxwell, we invariably produce more (total) entropy<sup>125</sup> amidst our attempts to sustain ourselves away from equilibrium by exerting energy, and therefore our acts do not violate the second law precisely because the *total* entropy of the universe is still increasing. Whilst I do not deem it prudent to elucidate the thermodynamic principles on display here in any more substantial detail, I do note that whilst Maxwell's initial sorting demon is not explicitly naturalizable without qualification, it is perhaps possible, as I shall elucidate, to conceive of a sorting demon that can employ some phenomenal apparatus to rearrange its structure in such a way that it may achieve greater access to useful energy and thereby act as a dissipative disequilibrium system working in a not dissimilar, but wholly naturalised, way to Maxwell's sorting demon. I now turn to provide an account of a Demon of this sort<sup>126</sup>.

#### 4.2.1 Phenomenal properties realizing Phenomenal Maxwellian Demons

Whilst Maxwell's dualistic commitments<sup>127</sup> perhaps precluded meaningful discussion of a truly naturalised perceptual sorting Demon, the above articulation of a demon equipped with certain phenomenal apparatus would have perhaps been favoured by Maxwell, for, at core, Maxwell was motivated by an attempt to delineate precisely how conscious minds act upon

---

<sup>125</sup> By 'total' entropy, I here mean the entropy of the universe as a whole. In this sense, a singular system within the universe may make itself more orderly as long as the total disorder of the universe increases as a result of this orderly behaviour. This is how dissipative disequilibrium systems are reconciled with the second law.

<sup>126</sup> Although I note that whilst Maxwell's Demon is directly sorting external objects, my Demon is individuating between its own phenomenal states and subsequently employing this in order to sort the external objects that these phenomenal states weakly represent.

<sup>127</sup> See Maxwell (1853), p. 117 and Maxwell (1862), p. 712.

physical processes. As Maxwell himself noted, the interaction between the conscious ‘soul’ or mind need not be one of a brute exchange of energy, but instead may be framed as a subtle force that guides the flow of energy:

*‘There is action and reaction between body and soul, but it is not of a kind in which energy passes from one to the other – as when a man pulls a trigger it is gunpower which projects the bullet, or when a points man shunts a train it is the rails that bear the truth. But the consciousness of our nature is not explained by finding out what it is not. It is well that it will go, and that we remain in possession, though we do not understand it’* (Maxwell 1826, p. 712).

The point here is that conscious causation may act not as the pistol projecting the bullet, but as the subtle, informational force that begins the causal chain of events that lead to the projection of said bullet. In this sense, Maxwell construes conscious causation as a ‘delicate force that initiates a larger process, like a pebble starting an avalanche’ (Stanley 2015, p. 469), such that ‘the temperature of B may be raised and that of A lowered without the expenditure of work, but only by intelligent action of a mere guiding agent (like a pointsman on a railway with perfectly acting switches who should send the express along one line and the goods along another)’ (Maxwell 1826, p. 469). The problem Maxwell faced, however, was that as a self-proclaimed natural dualist (see Maxwell 1853, p. 117), he persistently failed to reconcile the conscious mind with an adequate account of precisely how our mental states may interact with the physical to the extent that would be necessary to behave as this ‘delicate force’.

Conversely, I maintain that for a Russellian Micropsychist such reconciliation should not be overly troubling, for on this account physical entities simply *are* phenomenal. I therefore argue that whilst Maxwell's initial formulation of his sorting demon offered a neat means of integrating information into our understanding of thermodynamics, we may build upon Maxwell's ideas by thinking of them in terms of the commitments underpinning Russellian Micropsychism. If we do this, I maintain that it is entirely possible to conceive of a minimal, simplistic 'sorting' Demon that works purely phenomenally 'from the inside' to individuate between phenomenal types to sustain its access to experiential content, whilst 'from the outside' behaving like a dissipative disequilibrium system that is employing some unobserved mechanism to arrange its structure to provide greater access to the currents than run through it. I hold that this process would satisfy an account of a Naturalised Maxwellian Demon that works phenomenally to sustain a state of disequilibrium.

Construed as such, I note that this formulation is fundamentally distinct from Maxwell's sorting Demon but remains consistent with Maxwell's grander notion that consciousness may act to guide the flows of energy by use of, perhaps, some sort of broadly construed information (or as I shall argue phenomenal experiences). With this in mind, I note that whilst Maxwell's Demon was initially framed as a perfect sorting mechanism capable of tracking the movement and position of each individual particle, my sorting Demon is far clumsier and more simplistic, and works only to individuate between its base phenomenal state and any changes to this state. In this sense, whilst Maxwell's Demon was complex, my Demon is, at foundation, absolutely simple. It exists in its base structural state intrinsically phenomenally representing its own structure and, given a structural relation that is adequate to change the structure of this Demon, it is capable of individuating between the phenomenal state that occurred prior to the structural change and the phenomenal state that occurs as the

phenomenal representation of the new structure overlaps the phenomenal representation of the previous structure. I subsequently hold that the only real similarity between mine and Maxwell's formulations is that both Demons are, in theory, capable of employing certain broadly informational dispositions in order to sustain themselves away from thermodynamic equilibrium. In this respect I take it that if the necessary and sufficient condition for the occurrence of a Maxwellian Demon is that a given entity or system can employ some form of information (but not necessarily narrowly physical information) to slow the decline of either itself or its environment into a state of maximal entropy, then a system employing phenomenality in order to achieve this same effect would satisfy a broad account of what we mean by a Maxwell's Demon.

If we think of this in fundamental microscopic terms, we might construe this Demon as Busby & Howard (2017) do as a 'Quantum Demon', but caveat this with the contention that this 'Quantum Demon' works phenomenally. Busby & Howard elucidate the physical behaviour of such a Demon as follows:

*'Molecular species can be separated in energy terms through specific vibrational exchanges which have resonant absorption for certain molecular structures but not others. Thus, irradiating a mixture with infra-red radiation at a specific frequency which is absorbed by one molecule but not others immediately allows the resonant molecule to obtain extra energy and become hotter than its surroundings. This excess may be used by the chosen molecules to react with other 'cold' molecules in their environment and thus to grow, extend themselves etc' (p. 172).*

What this means is that if the universe fundamentally contains the potential to realise Phenomenal Demons (P-Demons), then certain instantiations of phenomenality may occupy the right kind of structure necessary to realise P-Demons and may very well then employ both an adequate source of energy and their disposition for phenomenal individuation in order to restructure themselves in the direction of less entropic states and thereby essentially maintain themselves by ‘eating their environment’ (Busby & Howard 2017, p. 172). The key point here is that the quantum demons work by measuring the position of cold molecules in the environment and then structuring themselves so as to provide access to these particles, and as the nature of this process of measurement remains open to interpretation, I hold that this measurement process may very well be purely phenomenal.

One way to make sense of how this sort of measurement may occur is to employ the notion of molecular energy ‘flows’<sup>128</sup>. On this account, as Elitzur (1994) elucidates, the Demon need not carry out complex measurements pertaining to the position of each individual particle, but instead every interaction with an energy flow conceivably constitutes a ‘primitive form of measurement’ (p. 444) because an interaction with one molecule ‘probably foretells a group of other molecules with the same momenta’ (p. 445). I take it that such an account is entirely consistent with my formation of Phenomenal Maxwellian Demons, for upon a structured subject interacting with an entity that is sufficient to change its structure, the structure of the subject changes to experientially represent the flow of this entity and, as a result, the phenomenality of the structured subject may be employed to arrange its structure as to provide the subject with greater access to more energy by restructuring itself such that it moves in the general direction of the flow of this particular entity. Such an account could

---

<sup>128</sup> By the term ‘energy flow’ I am referencing the direction of the energy that manifests as molecular entities move in unison.

then be employed to substantiate precisely how the ‘chosen molecules’ in Busby’s analysis access the colder molecules in their environment, for according to Elitzur these cold molecules are likely to flow together, and therefore an effective P-Demon may very well structure itself to provide greater access to several cold molecules by experientially measuring the flow of a small number of molecules.

As I have iterated, however, this process is largely clumsy and random, and therefore an interaction of this kind does not guarantee that the P-Demon will react to effectively restructure itself in a manner that may provide greater access to these energy flows, especially as it seems just as likely that a given structured subject will restructure itself to decrease its access to energy flows. I note, however, that a reaction of *some kind* is all that is necessary to ground P-Demons because it seems likely that effective instances of restructuring will be favoured by natural selection. In this sense, I note that P-Demons may very well work to effectively slow their decline into a state of maximal entropy by restructuring themselves as a result of the experiential ‘spikes’ they undergo, but these P-Demons remain absolutely primitive when compared to their Maxwellian forebears.

With all of this said, I note that whilst framing these types of behaviour in terms of phenomenality is relatively radical, any initial concerns might be curtailed because an account of precisely what drives certain systems towards self-sustaining non-equilibrium states is still a point of contention in contemporary physics (see Schrodinger 1944 or Walker 2017 for surveys of this problem). I therefore suggest that it remains entirely possible to explain this as Naturalised Demon-like behaviour and subsequently to think of phenomenal properties as responsible for sustaining certain systems in non-equilibrium states by virtue of grounding this Demon-like behaviour. I maintain that such an account need not commit us to

overdetermination concerns<sup>129</sup> and is no less naturalised or explanatorily powerful than an account predicated upon narrow physicalism (See England 2009 for such an account) because there is no discrepancy in the a posteriori commitments employed to substantiate such models and both retain the same predictive values. I therefore motivate this on the grounds that upon endorsing such an account we have explicitly explicated the potential behaviour that may manifest as a result of phenomenal properties' disposition to individuate, and thereby have gone some way to laying the foundation necessary to position phenomenality firmly into our naturalised model of the universe in a manner that explains what difference it makes to our understanding of the natural world without violating our a posteriori commitment to model formation.

In order to justify this *prima facie* radical call to construe the behaviour of phenomenal properties in terms of a revised account of Maxwell's (1867) sorting demon, I deem it necessary to first reiterate precisely what I take phenomenal properties to be and thereby reiterate the essential powers that occur as a result of their intrinsic qualities.

1: Bottom-level Russellian Micropsychist phenomenal properties are necessarily bottom-level, structured subjects of experience.

2: Bottom level phenomenal subjectivity is necessarily a causally significant powerful quality.

3: A bottom-level phenomenal subject is necessarily intrinsically, continuously experiential, and knows 'what it is like' to undergo experiential content.

---

<sup>129</sup> As iterated, this is because phenomenality is itself fundamentally causally relevant, but a more specific account why this is the case is forthcoming in the subsequent sections – see 3.4.3 specifically.

4: The extent of a bottom-level subject's experiential content is dependent upon the structures the bottom-level subject occupies.

5: If all of the above conditions hold, and they must, then, upon standing in the right relation to external structural properties, bottom level subjects of experience are capable of minimal, first-order phenomenal individuation as one structured state overlaps the next.

With this reiterated, we might simplify the core of the argument that underpins this section as follows:

If Phenomenal properties are powerful qualities, and

1. Phenomenal properties have the quality of being structured subjects of experience.
2. Structured subjects have the power to know what it is like to undergo content and are capable of first-order phenomenal individuation.

Then, given optimal conditions and realism about the laws of nature, this power will physically manifest as a behaviour that is physically explicable in terms of a Naturalised Maxwellian demon.

I take it that the core of the argument within this section is therefore relatively simple, and may even be simplified further: if phenomenal properties hold the quality of being a structured subject, and this quality entails the power to individuate between different phenomenal (and thereby structural) types, then structured subjects meet the necessary and sufficient conditions to behave like a phenomenal, naturalizable version of Maxwell's sorting demon<sup>130</sup>, when 1) appropriately integrated into a naturalised account of reality, 2)

---

<sup>130</sup> I reiterate that the account of Maxwell's demon I am offering is minimal and thereby wholly naturalizable. As has been established, it is fundamentally distinct in kind from the account initially offered by Maxwell and also avoids the traditional problems erected against Maxwell's Demon.




undergoing an interaction sufficient to cause a structural change, and 3) occupying a structure that does not preclude the possibility of restructuring to pursue more structural changes (i.e., occupying a structure replete with some mechanism necessary to restructure, or utilise free energy, as a result of phenomenal individuation). As shall be explored in the subsequent sections, when caveated as such it appears reasonable to infer that whilst phenomenal individuation might be a power associated with all instantiations of phenomenality if the conditions necessary to induce phenomenal overlap are present, the power to act like a P-Demon is only present upon a given instantiation of phenomenality satisfying each of the conditions outlined above.

Even with this caveat in place, I note that such a formulation is not without explanatory issue, because 1) it is not immediately clear that this behaviour is reconcilable with a universe that often seems to *not* structure itself to access more energy flows, 2) it is not clear precisely why phenomenal subjects would behave in this way, and 3) even if phenomenal individuation is granted as a potential, relational power of phenomenal properties, it is not clear that this necessarily enables structured subjects to act like a P-Demon. Whilst all three are troubling, I begin with the latter, for I take it that the integrity of the former now rests upon a sufficient account of precisely how structured subjects equipped with the capacity for phenomenal individuation results in the potential for P-Demons (if the term ‘P-Demons’ is employed in the liberal sense to denote any system that employs *experience* to as to sustain itself away from maximal thermodynamic equilibrium).

## 4.2.2 On the possibility of P-Demons

If the metaphysical picture I have painted thus far holds, the universe at base contains a multitude of unified and bounded structured subjects of experience, each undergoing their own intrinsic representational content that is derivative from the particular structure they occupy. Each such subject knows ‘what it is like’ to be in this base structural state, and, conceivably, it equally knows ‘what it is like’ to exist in the state that occurs as this base state is overlapped by a different structural state – in this sense, a structured subject can conceivably phenomenally *individuate* between the base state and the overlapping state in the moment of overlap. As a phenomenal overlap of this kind is itself caused by a change in the subject’s structure<sup>131</sup>, it is conceivably possible that whenever an external relation causes a change in a subject’s structure, a structured subject undergoes phenomenal overlap as the phenomenal representation of a previous structured state transitions into a representation of the new ‘post-collision’ structured state. If this is possible, then it seems equally possible that as structured subjects can conceivably phenomenally individuate between these two states, the subject could very well employ the experience of the overlap to structure itself in a manner that may lead to *more* instances of phenomenal overlap. We might elucidate this pictorially as follows:



*Figure 1: The potential physical consequence of phenomenal individuation*

---

<sup>131</sup> I.e., a change to the structure the subject is in.

Whilst such behaviour is difficult to adequately convey, the point is that if the smaller blue circles are taken to denote energetic structures with little to no potential to perform thermodynamic work, and the larger red circle is taken to denote a structured subject with some available free energy (i.e., energy that may be employed to perform thermodynamic work), then conceivably, if the dispositional makeup of the universe allows, the energy of the red circle can be sustained by employing experiences. And, arguably, this is achievable because all the structured subject need do is undergo instances of phenomenal overlap and employ its free energy to arrange itself such that it is directed towards the structure that caused the overlap. This need not manifest as entirely directed or purposeful behaviour, for all that is reasonably required is that structured subjects react when an overlap occurs. Many such subjects will likely react incorrectly to distance themselves from the structure, but equally some, by virtue of chance, will likely react appropriately and thereby may conceivably sustain themselves, at least for a while, away from equilibrium by employing experience to sustain their access to more overlapping states and thereby more potential (or free) energy. In this sense, *experience begets experience*. Crucially, the point to be laboured is that any energetic interaction between a structured subject and its environment is sufficient to cause at least some structural change and thereby some phenomenal overlap, and therefore it is not particularly difficult to conceive of how this could be applied to Elitzur's (1994) formation of a rudimentary and primitive form of measurement, for the phenomenal overlap is itself sufficient to perform a sort of very simplistic self-measurement that orientates the demon to one's immediate environment. All that is therefore required within an experiential model of Maxwellian Demons is the occupation of a structure that does not directly preclude the possibility of existing in a non-equilibrium state which grants some access to a mechanism capable of utilising free energy, i.e., some energy that is available to perform

thermodynamic work, and a relational interaction with an environment that is sufficient to cause some phenomenal overlap within a given structured subject.

As Busby & Howard (2017) elucidate, certain systems may gain access to free energy by absorbing radiation, or indeed by some other as yet unexplored means, but how this is achieved is relatively unimportant if we maintain that 1) it is conceivable that structured subjects have the capacity for phenomenal individuation when standing in certain relations to their environment, and 2) it is conceivable that such subjects should, at some point in the evolution of the universe, both have access to free energy and occupy a structure that would not directly preclude the possibility of employing this free energy to sustain a non-equilibrium state. I hold that neither of these inferences are beyond the scope of comprehension, and maintain therefore that as this is at least conceivable, then it is possible in this actual world, and if it possible in this actual world, then as this behaviour is clearly advantageous to survival, it is likely that this behaviour would have evolved in complexity and efficiency. Construed as such, the second issue highlighted above pertaining to the motivation of P-Demons is of relative unimportance, for regardless of whether a P-Demon is motivated to pursue overlapping states or not, as soon as there is a reaction to an overlapping state that is veridical and advantageous to the fitness of the system, the trait of moving towards more experiential states is likely to be selected for survival<sup>132</sup>.

With the speculative nature of some of these inferences set aside, it does at least appear to be the case that, if the second law of thermodynamics holds, the role of life (or consciousness) might be to avoid the decline into maximal entropy, as Pinker (2019) attests:

---

<sup>132</sup> See chapter 6 for an account of why this is likely to be the case.

*‘The second law defines the ultimate purpose of life, mind, and human striving: to deploy energy and information to fight back the tide of entropy and carve out refuges of beneficial order’ (p. 18).*

In this case, we might suggest that the cause behind these refuges of beneficial order is phenomenal in nature, for the move to avoid maximal entropy is ‘from the inside’ a move to sustain access to more structural changes and thereby more phenomenal content. Suffice to say that whilst this is speculative, construing experience as a cause of (at least some of) these refuges would arguably solve the problem of delineating precisely why and how information came to act upon physical processes and why and how some systems sustained themselves away from disequilibrium in the first place (i.e., the problem(s) of life<sup>133</sup>). With this final speculation in place, I note that the first issue facing an adequate account of P-Demons remains, for if phenomenal properties are ubiquitous, one must explicitly account for precisely why P-Demons are not. I now turn to address this concern by offering a more detailed elucidation of the ontology of causation I am offering, which will itself provide an account of causal blockers that may find a place for both phenomenal individuation and P-Demons (and thereby explain away this aforementioned problem).

#### 4.2.3 A place for Phenomenal Individuation and P-Demons

*‘Even if there is only one possible unified theory, it is just a set of rules and equations. What is it that breathes fire into the equations and makes a universe for them to describe? The*

---

<sup>133</sup> Walker & Davies (2016) describe the problem of life as a ‘hard’ problem and suggest that it is of comparable import to the ‘hard’ problem of consciousness. If this is right, a theory of P-demons might lead us to avoid some of the explanatory burdens afflicting mainstream physicalism.

*usual approach of science of constructing a mathematical model cannot answer the questions concerning why there should be a universe for the model to describe. Why does the universe go to all the bother of existing?' (Hawking 1989, p. 174)*

As Hawking intimates, physics provides us with a sufficient mathematical model for the universe but does not provide us with a complete account of 1) whether the universe has a nature that is not captured within the rules and equations or 2) the cause of (or motivation for) the behaviours documented in these mathematical models. With my advocacy of Russellian Micropsychism in mind, we might suggest that the first of these issues has been solved by virtue of the commitment to the notion that the universe is intrinsically phenomenally experiential at the bottom-level and that all facts supervene on the bottom-level instantiations of these powerful phenomenal-physical qualities. This does not in itself, however, provide an account of precisely why the universe is fundamentally motivated to behave in the manner that is captured by the mathematical models employed in physics, and so the second issue arguably remains open.

Given the accounts of phenomenal individuation and P-Demons I have offered, it may be tempting to speculate about the intrinsic motivations of the universe by maintaining that the disposition to phenomenally individuate manifests at the very bottom-level to necessitate a myriad of rudimentary P-Demons that guide the universe towards more experiential states. As tempting as a teleological account such as this might be, I hold that it is simply not possible to conceive of a way in which we might make sense of phenomenal individuation, or worse P-Demons, manifesting at the bottom-level if the bottom-level is itself fundamentally simple and is curtailed by various natural laws. To suggest that powerful phenomenal-physical qualities are disposed, at the bottom-level, to employ their capacity for phenomenal

individuation to create the experiential universe as we know it would therefore require a revisionary account of physics that is simply entirely beyond the scope of what is achievable within the confines of this thesis<sup>134</sup>. Instead, I suggest that we do *not* revise our physics to accommodate bottom-level P-Demons and instead simply suggest that P-Demons only occur at higher-levels of the universe, such that if phenomenal individuation is perhaps conceivable as a disposition of instantiations of phenomenality that exist at the bottom-level, P-Demons are conceived of exclusively as a disposition of instantiations of phenomenality that exist at higher levels (i.e., states of the universe in which phenomenality acquires parts/complex structure). At core, we might think of this as simply two ways of conceiving of the same disposition (phenomenal individuation) whilst remaining true to the physics by suggesting that our current understanding of the universe precludes a physical manifestation of this power at the bottom-level. In this respect, we might think of phenomenal individuation as a seed patiently waiting for the sunlight so that it may flourish. Like the seed, once phenomenal individuation is standing in the right reciprocal relation to its environment, it *will* flourish, and it *will* gain tangible causal traction on reality. The causal traction it gains is, as I have argued, best thought of in terms of ‘P-Demons’<sup>135</sup>.

---

<sup>134</sup> I do note that one way of perhaps achieving this account would be to posit bottom-level phenomenality as responsible for collapsing the wave function. This would reconcile micropsychism with the von Neumann-Wigner interpretation of quantum mechanics by suggesting that it is bottom-level phenomenal consciousness that provides the sort of observation necessary in order to move the wave function from a superposition of several eigenstates to a single eigenstate. For my purposes, I simply note that this is one way that a theory of this kind might be postulated. Alternatively, we might adopt Peirce’s (1857-66) account of evolved natural laws and posit that the specific natural laws we observe are simply those laws that have been evolutionarily optimised to perpetuate the experience of the cosmos. I do not, however, address either of these theories in any more depth here, because, as intimated, I take it that a revisionary account of our understanding of physics is, in this instance, beyond the scope of this particular thesis.

<sup>135</sup> I note here that if P-Demons are possible (as I will argue they are), we might employ them to solve what I reference as the ‘second-order combination problems’ (see chapter 6). I also note, however, that prior to this we need to provide an account of what made P-Demons possible in the first place: we need a solution to the subject-summing problem that might account for how a myriad of formed phenomenality might combine to form a P-Demon so that phenomenal individuation might reasonably manifest as I describe. A solution to this subject-summing problem shall be advanced in chapter 5.

In this respect, whilst we might hold that all higher-level properties supervene on the various bottom-level ‘aspects’ of the fundamental identity of the powerful quality of phenomenal-physicality, this is not a claim that this fundamental powerful quality is not continuous with physics. The natural laws as described in the standard model hold, and all the dispositions described in the standard model manifest precisely as described in the standard model. In this sense, the mathematical model for the universe is a description of the ways in which the dispositions of powerful phenomenal-physical qualities manifest in the world given the way the world is (i.e., given the various starting conditions evident in this universe and given the various fundamental forces and natural laws at play). Therefore, the fundamental powerful-quality I describe behaves precisely as the mathematical model for the universe suggests, and it cannot employ any of its dispositions to violate natural laws or perform the impossible.

How, then, might we find a place for phenomenal individuation and P-Demons that is entirely consistent with the mathematical models described by physics? In addressing this, I suggest that we must distinguish between a disposition and its causal significance. On this distinction, we might hold that all ‘aspects’ of phenomenal-physicality are disposed to phenomenally individuate when they are standing in a relation to their environment that is sufficient to cause a change to their self-representation of their form/structure, but that this disposition does not in itself manifest as a P-Demon unless the aspect of phenomenal-physicality that is instantiated is capable of physically reacting to the individuation that has occurred to sustain access to experiential overlap. If it is not, the subject might experience phenomenal overlap, and individuate a difference in phenomenal content, but this overlap would simply wash over the subject without any subsequent physical manifestation, or physically realised causal consequence of this overlap, that would point to this having occurred. In this respect, we might say that whenever quarks, for example, undergo colour change, they might



phenomenally individuate between the previous experiential state and the overlapping experiential state<sup>136</sup>, but there is no space for there to be physical evidence of this capacity for individuation at the bottom-level because for there to be we would have to revise our understanding of how the strong nuclear force works<sup>137</sup>. So, we might say that the physical consequences of phenomenal individuation only manifest when we reach higher states of the universe - states that might reasonably allow for freely morphing systems that exist far from equilibrium (i.e., the states necessary for P-Demons)<sup>138</sup>, and I now turn to expand on this account by offering a theory of dispositions and their manifestations framed in terms of reciprocal relations and causal blockers.

To frame this account, I begin by highlighting that whilst bottom-level instantiations of formed phenomenality are bundles of powerful qualities that manifest as various ‘aspects’ of phenomenal-physicality ubiquitously at the bottom-level, much in the same way we would not expect a spherical ball to roll unless placed on an incline, these powerful-qualities are tied

---

<sup>136</sup> This is possible because quarks change their qualities (i.e., mass) when they change colour/flavour, and so, as this is a process that is thought to occur continuously at the bottom-level, we might say that quarks conceivably continuously phenomenally individuate between these changes in their form/structure.

<sup>137</sup> I note that the strong nuclear force might very well be, in some sense, ‘guided’ by phenomenality. However, as noted previously, I take it that there is no room for a detailed revisionary account of this kind within the confines of this thesis.

<sup>138</sup> By this I mean the conditions necessary for P-Demons to function. These conditions include the absence of a force that would preclude the manifestation of this behaviour, a sufficient source of available free energy necessary to restructure and sustain a non-equilibrium state (and a mechanism that would facilitate the use of this free energy for this purpose), and sufficient interactions with an environment necessary to incite structural changes in phenomenality. We might describe these conditions with reference to a state in which self-organising dissipative structures (Prigogine 1993, Bejan & Lorente 2011) first manifested, for a structure of this kind would satisfy these conditions in totality. Whilst I note that there is a dispute in the scientific community pertaining to precisely when this state occurred (see Jia et al 2019 or Michaelian 2017), there is little to no dispute that such a state *did* occur *at some point*. For my purposes I take it that delineating the precise point in which this state occurred is of less importance than acknowledging the uncontroversial claim that, in fact, a state (i.e. a state necessary to allow for the possibility of self-organising dissipative structures) *did* occur *at some point* in the history of the universe, and I argue that this concession is of such importance simply because such a state is all that is required in order to point to at least one period in which the reciprocal relations at play in the universe did not preclude the onset of P-Demons. With this established, I will not devote any more time to arguing for the possibility that a state in which the right reciprocal relations were in place for the manifestation of P-Demons. I take it that as it is uncontroversial to describe the origins of life in terms of self-organising dissipative structures (see Prigogine 1993, Kondepudi et al 2017, Michaelian 2017 amongst various others), it is not unreasonable to maintain that our existence as living systems points to the reality that such a state did occur.

to the dispositional makeup of the universe at a given time. It may be the case that, as Heil (2012) notes, certain powers associated with this quality are occasionally ‘blocked’ in the same way that the intake of a poison may be blocked by an antidote. This is not to say, however, that the poison is removed entirely, it is simply an attestation that the disposition of the poison is behaving exactly as it ought to, given the reciprocal relations at play. A similar analogy might be employed to explain the nature of micro-level phenomenal properties. Whilst they continuously physically manifest and realise the dispositions that occur at the bottom level, the dispositions we might expect them to manifest are dependent upon the manifestation of reciprocal relations that do not act as blockers to the disposition they ought to, or can potentially, manifest. In this sense, if a given disposition does not manifest, it does not imply that powerful phenomenal qualities are not manifesting, it simply implies that, given the state of the relations at play, they are manifesting precisely as would be expected.

Heil (2012) makes this point clearer by employing Descartes’ analogy of a malfunctioning clock that has not been sufficiently lubricated. Whilst the malfunctioning mechanism might lead us to attest that the clock is not working as it ought to, in fact the clock is behaving precisely as it ought to, given its dispositional makeup. In this analogy, its dispositional makeup is derived from the reciprocal relations at play, and, in this case, these relations are missing a key component: lubricant. As Heil notes, if the lubricant were added to the clock’s dispositional makeup, the manifestation of its behaviour would be changed. Similarly, whilst structured phenomenality may be said to exist at the bottom-level and realise all fundamental dispositions, it is only upon adding the right hypothetical ‘lubricant’ to this phenomenality that we might coherently be able to refer to the behaviour of this phenomenality as the manifestation of a P-Demon. We might therefore say that powerful phenomenal qualities hold the power to realise all of the dispositions observed in physics by virtue of being

fundamentally structured (i.e., being fundamentally physical and thereby behaving exactly as non-phenomenal physical ultimates might at the bottom-level), but also realise the power to cause some systems to behave as if guided by P-Demons, unless otherwise constrained.

To disambiguate this sentiment, we might elucidate the core of this account by way of an analogy. Imagine a ball with the property of sphericity. This ball, by virtue of being spherical, may potentially manifest a number of distinct dispositions. It may, for example, fall through a round hole, it may roll, it may look spherical, it may imbed its sphericity unto space-time et cetera. In each case, this sphericity manifests as a distinct disposition when interacting with certain reciprocal partners, such that the manifestations of the disposition of sphericity is dependent upon its relation to its environment. Similarly, we might say that as physically formed subjects of experience exist at the bottom-level of reality, they are the qualitative undercurrent of every possible state the universe may exist in, but their dispositions remain tied to the possible reciprocal relations at play in the universe at a given time. In this sense, all I am doing by integrating the concept of powerful phenomenal qualities into our worldview is 1) reducing the dispositional makeup of the universe to fundamental physical-phenomenality, and 2) imbuing our ontological landscape with a *potential* that would not have occurred in the absence of phenomenal qualities – phenomenal individuation. From here, I am employing an adaptation of the Heilian (2012) account of powerful qualities and reciprocal relations in an attempt to position this potential into our worldview in a manner that, for the most part, innocuously reconciles physical-phenomenality with the standard model, but also, given the right reciprocal relations, provides a means of displaying precisely how a specific aspect of phenomenal-physicality may behave like a naturalised Maxwell's Demon and thereby may render the role phenomenality plays in the natural world explicit in a manner that is consistent with liberally naturalised metaphysics.

I am subsequently simply suggesting that if physical ultimates are fundamentally phenomenal, and this phenomenality is fundamentally a property that subsumes several bottom-level powerful qualities, then the powers we expect to be associated with phenomenality will manifest in the universe *if* the universe's fundamental forces and powers occupy a state in which the reciprocal relations manifest that are necessary for the manifestation of these powers. Therefore, we might return to the sentiment that the identity of phenomenal-physicality is instantiated in a number of 'aspects', and each of these aspects instantiate various qualities and dispositions whilst sharing the category of being phenomenal-physicality. Amongst these 'aspects' of phenomenal-physicality, there is a state of being 'X' that undergoes phenomenal overlap as its physicality is changed by virtue of interacting with its environment, and this state of being is disposed to phenomenally individuate between the pre-collision state and the post-collision in the moment in which the phenomenal overlap occurs. Similarly, there is another 'aspect' of phenomenal-physicality occupying a state of being 'Y' that satisfies 'X', but also holds various other properties that dispose it to employ its phenomenality to react appropriately when the overlap occurs and restructure itself in the direction of more overlap. The former instantiation is a crude articulation of phenomenal individuation and the latter is a crude articulation of a P-Demon. In both cases, the manifestation of the dispositions is tied to the reciprocal relations at play in the universe, such that if they *do* provide a means in which this causal power may manifest as expected, I hold that this causal power will behave like a P-Demon that is equally naturalised given the dispositional makeup of the universe at the time in which the demon manifests. Conversely, however, if the reciprocal relations do *not* provide a means in which this causal power may manifest, this does not imply that phenomenal properties are causally inert, it is

instead just to say that the powers associated with structured phenomenal qualities are manifesting exactly as they ought to, given the reciprocal relations at play<sup>139</sup>.

In this sense, the behaviour of structured subjects that act like P-Demons is explicit – it is a reactant restructuring towards more experience – but this behaviour remains wholly consistent with the standard model because any observational evidence to the contrary (i.e., any evidence that matter does not seem to act in this directed manner) is explained away as the by-product of a dispositional state of the universe in which this particular directed behaviour is blocked, be that partially or entirely. Further still, there can be no overdetermination concerns because there is a brute identity between phenomenality and physicality which entails that a disposition associated with a bottom-level broad physical quality is in itself a disposition associated with a phenomenal quality, such that phenomenal individuation has a place and role in the causally closed physical universe because 1) phenomenality potentially just *is* those physical qualities associated with certain disequilibrium and brain-states, and 2) there is only one fundamental cause for all of the dispositions observed in physics: the powerful quality of phenomenal-physicality (or, alternatively, the property of structured phenomenality that subsumes all fundamental powerful qualities)<sup>140</sup>. In this respect, we can outline this account as follows:

---

<sup>139</sup> Further, this is not to suggest that there is not some directedness entailed by the powerful quality of physical-phenomenality, for in line with Molnar (2003, p. 60) it seems perfectly coherent to posit that directness is an ‘essential feature’ of powers (and thereby also powerful qualities). As Molnar (2003) elucidates, we might describe the directedness of gravitational mass in terms of an orientation or purposiveness towards the exertion of gravitational force. Similarly, we might describe sphericity as being directed towards rolling, or perhaps ease of motion. Either way, the point here is that powers are directed towards their manifestation unless otherwise constrained, and so, in the case of structured subjects, we might ground the behaviour of Phenomenal Maxwellian Demons by stating that structured subjects are directed towards the attainment of more experiential states, unless otherwise constrained. In this sense, I hold that it is *prima facie* reasonable to infer that subjects are motivated by experience and this motivation does not necessarily create any contradictions with the natural sciences as long as we uphold the Heilian (2012) notions of blockers/reciprocal relations as partial realisers of potential powers

<sup>140</sup> As Goff (2019) notes ‘It is generally held that over-determination (there being two distinct sufficient causes for a single effect) is innocuous in cases in which one of the two causes is constituted by the other (Bennett

- 1) Structured phenomenality is disposed to all the manifestations observed in physics.
- 2) Structured phenomenality is disposed to at least one power not directly observed in physics— the capacity to phenomenally individuate.
- 3) Any given power associated with the quality of structured phenomenality may be blocked if the dispositional makeup of the universe is not conducive to the manifestation of this power.

I thereby reject the claim that there is no place whatsoever for additional accounts of the causal powers evident in the universe and, as I have done throughout, substantiate this on the grounds that the current physicalist account of physical dispositions is limited in so far as our best epistemic tools can only allow us to reasonably make inferences pertaining to the structural behaviours observed in the empirical sciences, not necessarily the intrinsic cause or nature of these behaviours. Therefore, as long as a worldview that provides a different ontological account of a given disposition is empirically equivalent to the more scientifically ‘standard’ account of these powers, an analysis of these two accounts should proceed by way of analysis of their respective explanatory power. In this case, powerful phenomenal-physical qualities retain consistency with the empirical sciences entirely (and avoid the ‘hard’ problem) whilst also offering an explicit account of what difference phenomenality might make to our understanding of some of the powers at play in the natural world. With this said, I now turn to compound the case for the naturalisation of P-Demons by offering an account of Bejan’s Constructal Law that might be employed to reconcile the phenomenal behaviour of

---

2003; Goff 2017)’ (p. 103). In this case, as phenomenal-physicality is the fundamental constituent of the universe, it constituted all other causes and thereby avoids overdetermination concerns.

certain structured subjects ‘from the inside’ with an account of how this Demon-like behaviour might look ‘from the outside’.

#### 4.2.4 The Constructal Law ‘from the inside’

If the final section of chapter 3 was concerned with broadly naturalising structured phenomenality by way of placing this quality as the ground for the dispositions observed in physics, this section is concerned with naturalising one specific power that occurs by virtue of the quality of structured phenomenal experience – that is, the power to phenomenally individuate between representational content. As iterated, this power may conceivably be exercised to ground P-Demons, which may reasonably be described as a reference to any instance in which a naturalised system, or structure, sustains itself away from thermodynamic equilibrium by way of employing phenomenal experience. Construed as such, whilst it seems relatively simple to absorb the Russellian-Eddington account of panpsychism to show how the dispositions observed in physics are broadly realised by structured phenomenality, the problem of explaining how Demon-like behaviour can be naturalised must be supplemented by an account of a specific physical phenomenon that may reasonably be said to capture this behaviour ‘from the outside’. In this respect, because 1) P-Demons seem capable of employing experience to access more experiential content and 2) more experiential content is gleaned by way of increasingly fluctuating structural states (and thereby increased energetic states), we require a physical phenomenon that seems to encapsulate a universal time-direction towards more energetic, or more structured, states such that the flow of energy designs itself to provide access to more energy (and thereby ‘from the inside’ also more experience). This, at core, is Bejan’s (1997-16) widely championed ‘Constructal Law’.

This law, as Bejan (2005) elucidates, describes the occurrence of designs in nature, and predicts the time-direction of the evolution of such designs. In this sense, the law may be articulated as follows: ‘for a flow system to persist in time it must evolve in such a way that it provides easier access to its currents. This is the law of configuration generation, or the law of design’ (Bejan 2005, p.1677). This amounts to the observation that for a finite system to maintain itself away from equilibrium (i.e., to persist in time/live), it must configure itself, or to use a term often ascribed to Bejan, ‘morph freely’<sup>141</sup>, so that it provides greater and greater access to the currents that run through it (Bejan 2010, p. 1335). Construed as such, this law is not an elucidation of the relatively benign inference that for living things to remain living they must evolve such that they have increasing access to free energy. Instead, as Bejan (2010) himself articulates, at core the Constructal Law captures the observation that ‘nature is configured to flow and move as a conglomerate of ‘engine and brake designs’’ that move towards states in which greater access to energy flows is achieved (p. 1335). Evolution subsequently acts upon these designs to favour the generation of configurations that provide greater access, but it is the direction of this phenomenon that is of import, for it is this that, as Bejan (2010) elucidates, points to the ubiquity of ‘*the generation of configuration*, or the generation of ‘design’ in nature’ (p. 1340). At foundation, then, the Constructal Law is a commentary on the time-direction of the universe.

To articulate this, let us return to the analogy of the set of ‘infinity’ dice outlined within the introduction of this piece. As before, imagine that these dice act as the ontological foundation of the universe such that the way they roll reflects the possible states the universe may occupy. Now, imagine that there are two possible directions these dice may flow, or roll, in.

---

<sup>141</sup> See Lorenzini et al (2011, p. 196-120) for uses of the term in relation to Bejan’s Constructal Law.



Either they may roll towards a state in which they flow and roll less and less easily, or they may roll towards a configuration in which they may roll (flow) more easily<sup>142</sup>. The ubiquity of this latter direction is what the Constructal Law captures: the observed phenomenon that the time-direction of the universe is a movement towards more flow. Bejan describes this as follows:

*'All thermodynamic systems in nature are flow systems (i.e., live, non-equilibrium systems), and they all have configuration. If they do not have it, then they acquire it, in time. The generation of configuration is ubiquitous, like other phenomena covered by other 'laws' in physics. Biological systems are configured. Geophysical systems are configured. Engineering and societal systems are configured. The configuration phenomenon unites the animate with the inanimate. All the other phenomena of physics (i.e., of 'everything') have this unifying power.'* (p. 1336).

So, if Bejan is right, we have a unifying power that manifests as the construction of freely morphing designs in nature that evolve to provide greater access to the energy flows required to continuously 'freely morph' into more optimised designs/configurations. If this is so, then the question invariably becomes: what causes this phenomenon? What does evolution act upon? And why, precisely, is the universe configured to work in this manner? Bejan (2012) himself seems to have grappled with these fundamental questions upon suggesting that the phenomena of design in nature 'raises the question: How come? What causes the Constructal Law?' (p. 14). The short answer, according to Bejan (2012), is that, simply, 'we do not know'

---

<sup>142</sup> I note that this is not to say that the Constructal Law is necessarily at odds with second law, especially when one considers that easier, or more, flow always results in an expediated production of entropy. In this respect, this account of the Constructal Law might be thought of as another mode of presentation for this same phenomenon, although I note that the the Constructal Law does seem to capture something about the rate and design of the production of entropy that is not captured by the second law.

(p. 14-15). The brevity of this response should not be downplayed, for it is difficult, if not impossible, to delineate precisely why certain macro-level systems freely morph in the manner that they do. Or, more simply, it is difficult to delineate precisely what causes the universe to move in the direction that it does. If this remains an unknown, then I posit that there is room for this law to be reconciled with phenomenal properties in a manner that would make sense of the account of P-Demons I offered in the previous section.

On such an account, one could conceive perhaps of how an instance of structured phenomenality should ‘from the inside’ undergo experiential overlaps upon standing in the right relation to external structures and subsequently, clumsily, reconfigure itself in the direction of more experience as a consequence, whilst ‘from the outside’ this behaviour presents as a ‘freely morphing’ design configuration that moves in the direction of more flow and may subsequently, potentially, evolve to provide greater and greater access to the currents that run through it. In this case, these currents conceivably precipitate structural changes in the structured phenomenality that grounds Demon-like behaviour, and therefore ‘from the inside’ the Demon is reconfiguring because of experiential overlaps whilst ‘from the outside’ this reconfiguration is observed as behaviour that appears as if ‘nature is configured to flow and move as a conglomerate of ‘engine and brake’ designs’ (Bejan 2010, p. 1335). The point here is not that P-Demons are *necessarily* responsible for the Constructal Law when the law is described as a broad observation about the time-direction of evolution, configuration, or generation, of design, it is instead simply the observation that as 1) the fundamental cause of the observations underpinning the Constructal Law remains unknown, and 2) it is conceivable that phenomenal Demon-like behaviour should occur if the dispositional makeup of the universe allows, then once this has occurred the cause of the direction of freely morphing designs in nature, *in certain systems*, may very well be

phenomenal and *this would not violate the observations underpinning the Constructal Law*.

In this sense, as the most interesting aspect of the Construal Law seems to be the inference that nature is configured to design itself such that it provides access to more energetic states, I am noting that whilst this time-direction could be conceived ‘from the outside’ as a movement towards more energy flows, ‘from the inside’ this could be conceived as a movement towards more experiential states. This is because, if my account of structured phenomenality is correct, increasing energy flows produce structural changes that result in increasing experiential states, and so in accessing more experiential content, a given system is ‘from the outside’ accessing more energy flows. Construed as such, we might say that there are two modes of presentation for the Construal Law. From the outside, this law manifests as a time-direction towards more energetic states, whilst from the inside the law manifests as a time-direction towards more experiential states. Further, as the mechanism that grounds the initial configuration of certain systems remains ambiguous, it seems entirely possible that this is grounded, at least in certain systems, in structured phenomenality, and doing so would provide an account of Demon-like behaviour that is consistent with the observed phenomenon of design in nature.

I note that such an account would not be inconsistent with Bejan’s (2012) conception of ‘design’, for all his account entails is a rejection of any notion of an external ‘designer’, and a commitment to the notion that nature designs itself. On the Russellian-Eddington account of panpsychism I am offering, nature does design itself, but it does this because nature fundamentally *is* structured phenomenality, and structured phenomenality grounds the potential for a system to employ experience to move towards more experiential content and thereby sustain itself away from thermodynamic equilibrium. In this sense, this is not an account of ‘design’ grounded in any form of intelligence or a mystical vitalistic life-force, it

is instead energy flows guided by experience. It is an account of the time-direction of freely morphing flow systems towards more experience, and this behaviour, as I have attempted to show, would manifest precisely as the Constructal Law predicts, 'from the outside'. The observed tendency for systems to maximise flow is therefore construed 'from the inside' as a tendency for systems to maximise experience and, arguably, conceiving of phenomenal causation in this manner would be entirely consistent with Maxwell's motivation to explain mental causation as a 'delicate guiding force'. Indeed, in this model, this is precisely the role phenomenality might play, and precisely the role that is easily reconciled with the formation of P-Demons offered previously.

In this respect, I take it that reconciling an account of Maxwellian Demons with the Constructal Law should not pose any significant issue, especially as Bejan (2014) himself describes the Constructal Law in terms of Maxwellian demons within an article entitled 'Maxwell's Demons everywhere: Evolving Design as the arrow of time', in which he articulates the law as follows:

*Science holds that the arrow of time in nature is imprinted on one-way (irreversible) phenomena, and is accounted for by the second law of thermodynamics. Here I show that the arrow of time is painted much more visibly on another self-standing phenomenon: the occurrence and change (evolution in time) of flow organization throughout nature, animate and inanimate. This other time arrow has been present in science but not recognized as such since the birth of thermodynamics. It is Maxwell's demon. Translated in macroscopic terms, this is the physics of the phenomenon of design, which is the universal natural tendency of flow systems to evolve into configurations that provide progressively greater access over time, and is summarized as the constructal law of design and evolution in nature. (p. 1)*

Bejan's (2014) aim in this passage is therefore to construe Maxwell's Demon macroscopically to 'imagine "a being" that can follow the flow of heat and divert some of it to flow through a contrivance—a design, or machine—that produces power, mechanical or electrical', and subsequently employ this idea to show 'this happens everywhere in nature, from the whole earth as a heat engine, to every animal as a vehicle with its own motor' (Bejan 2014, p. 1), and as such my articulation of P-Demons is seemingly entirely consistent with the Constructal Law because this phenomenon describes 'from the outside' what I take certain structures, i.e., those structures capable of realising P-Demons, to be doing 'on the inside'.

If this is so, structured phenomenality is conceivably immediately integrated into our naturalistic account of the time-direction of the universe and therefore may conceivably be subject to evolution by natural selection, for it seems that only those P-Demons that can most effectively structure themselves to provide greater and greater access to experience (or energy flows when construed 'from the outside') are likely to persist in time. In this sense, I suggest that P-Demons will manifest wherever the universe permits, and these P-Demons are fundamentally caused by structured phenomenal subjects standing in the right reciprocal relation to their environment. As I cover in chapter 6, whilst these P-Demons may initially be posited as sluggish, imprecise phenomenal sorting mechanisms that dictate, in certain systems, the flow of matter towards more energetic (and thereby experiential) states, it is likely that evolution would act upon this underlying mechanism by selecting (and favouring) the most efficient manifestations of this behaviour. We might even say that we have a multitude of hierarchies of structured phenomenal subjects running through reality, and the extent to which these subjects manifest as effective P-Demons may therefore be said to

produce different structures that vary dependent upon where this Demon manifests within the evolutionary hierarchy. In essence, then, as P-Demons are necessarily structured, the extent and complexity of the structure of a subject is likely to be a marker for both the extent of the phenomenal subject's causal influence and experiences *and* the position in the evolutionary hierarchy the subject occupies. In this sense, the more rudimentary and simplistic the structure, the more simplistic and rudimentary both the subject and its powers, and thereby the lower down the evolutionary hierarchy this subject is likely to manifest.

With this in mind, I now turn to substantially expand upon these inferences by 1) outlining an account of how we might solve the standard combination problem(s) in a way that might leave the universe in the necessary 'higher-order' states to enable the causal expression of phenomenal individuation (and thereby realise P-Demons), and 2) offering an account of the Constructal evolution of the universe that could incorporate these P-Demons to offer both a coherent solution to what I term the 'second-order' combination problems for panpsychism *and* unravel the most fundamental of all mysteries: *what are we?*

## Chapter 5

### Addressing the ‘first-order’ combination problems

According to Bejan’s (1997-2016) Constructal Law, our existence at this moment is one frame in a very long ‘movie’ of design evolution. The reason we experience this design as we do, as I shall argue in what remains of this thesis, is because we are the macro-level conscious manifestation of the design of increasingly effective, and evolutionarily optimised, P-Demons. However, as was outlined in the various articulations of the combination problem addressed in the first chapter, such an account faces the problem of reconciling the simplistic qualities and powers entailed by P-Demons with an account of precisely how we experience reality with the degree of unified complexity we do. Whilst it seems intuitive to instinctively turn to evolutionary theory in an attempt to delineate an elegantly simple solution to this problem (and as I shall argue our instincts might very well be right), our instinct, in this case, must be tempered, for whilst such a move appears *in principle* capable of accounting for our complexity, a sufficient solution to this problem simply must first encompass a detailed account of precisely how P-Demons came to exist in the first place, and must therefore address how numerous ‘simplistic’ unified and bounded micro-level instantiations of formed phenomenality combined to form a unified and bounded macro-level structured phenomenality (if macro-level structured phenomenality is here taken to refer to phenomenality that has various micro-level parts). It appears therefore that any appeal to evolution as a sufficient solution to this problem must first address the combination problem afflicting panpsychism by providing a detailed account of precisely how subjects ground subjects.

With these problems highlighted, I suggest that we might combine both of these issues by describing the fundamental nature of the problem facing micropsychist accounts of panpsychism as follows: if a macro-level subject of experience is constituted from micro-level subjects in the sense that micro-level subjects are phenomenal-physical parts of the macro-level subject, precisely how do the innumerable individually unified and bounded simplistic and sparse phenomenal parts of this whole combine to form just one unified and bounded complex, rich, phenomenal, structured whole that has an awareness of its existence as a unified conscious subject<sup>143</sup>? To adequately address this, I suggest splitting the problem into two categories of varying priority. The first category shall be tackled in this chapter and contains, what I call, ‘first-order combination problems’: the standard problem of subject-summing (which, because of my articulation of the relationship between subjects and their content advanced in chapter 2, also by extension entails an account of content-summing) and the palette problem and the ‘inverted’ combination problem I outlined in chapter 1<sup>144</sup>. The second category, which shall be tackled in chapter 6, contains, what I call, ‘second-order combination problems’: the structural mismatch problem and the problem of accounting for precisely how a basic consciousness with simplistic qualities and powers transitions to the sort of complex structured consciousness that achieves the conscious awareness of itself necessary to make inferences about its constituent parts (i.e., human consciousness). The call to split these combinations problems is not an attestation that these issues are not equally

---

<sup>143</sup> I take it that this articulation contains the various aspects of the combination problem as described in Chalmers (2016). It contains the problems of 1) how do microexperiences come together to form a unified consciousness? 2) How do microexperiences come together to form a bounded consciousness? 3) How do microexperiences come together to form an awareness. It also less explicitly contains 1) the palette problem: how do a relatively simple and sparse accumulation of bottom-level experiences produce rich and varied macro-level experiences? and 2) the structural mis-match problem: how does the rich structure of consciousness correspond to the structures of the brain?

<sup>144</sup> I take it that a solution to one of these problems will entail a solution to the other because an account of how or why phenomenality combines will necessitate an equal account of why or how other instances of phenomenality do not combine.



severe or important<sup>145</sup>; instead, I simply hold that the first-order problems ought to be solved before we begin to address the second-order problems because a solution to the former set of problems lays the foundation for a solution to the latter set. With this in mind, in this chapter I offer an account of how we might go about solving the first-order problems and suggest that, if these are solved, we are left with macro-level instantiations of structured phenomenality (i.e., a unified and bounded phenomenality with parts) which may, in certain iterations, exist as a mode of being that is disposed to behave like a P-Demon. With this established, I suggest that the remaining second-order combination problems might be solved by integrating these instantiations of P-Demons within Bejan's Constructal Law to provide an account of how these demons might transition from relatively simplistic instantiations of consciousness to complex conscious subjects by organising themselves such that they have increasing access to experiences. In this sense, I propose a two factored approach to the combination problems for panpsychism: a solution to the first-order problems that lead us to P-Demons (chapter 5) and an account of the evolution of P-Demons that leads us to human consciousness (chapter 6).

## **5.1 Towards a solution to the first-order combination problems**

As iterated in section 1.3, the combination problem is perhaps the most fundamental problem afflicting contemporary panpsychism. The issue, at core, is one of reconciling precisely how numerous unitarily bounded experiences combine to form a singular macro-level unitarily bounded experience of the type we seem to have as human subjects. Given my articulation of phenomenal properties in chapter 2, we might formulate this problem in terms of structured

---

<sup>145</sup> I maintain that each of these problems must be solved in order to provide an adequate panpsychist theory of consciousness, and so one is no more important than the other. Contrarily, see Coleman (2014, p. 29) for an articulation of one account in which subject summing is described as the 'real' combination problem and is thereby posited, in some sense, as more severe, or important, than what I describe as the second-order problems.

phenomenality by thinking of it in terms of an account that might explain how multiple bottom-level instantiations of structured phenomenality, or formed subjects of experience, come together to constitute a singular formed subject of experience. In this respect, the subject-summing problem and the combination problem are two ways of describing the same issue: how do simple unitarily bounded micro-level subjects of experience combine to form a singular unitarily bounded macro-level subject of experience? As noted, this problem does not in itself contain the problem of how awareness occurs, how complex phenomenally rich experiences form, or how complex structures equate to complex consciousness. In this sense, this problem exclusively addresses how unified and bounded subjects ground other unified and bounded subjects. More precisely we might suggest that if, as I argued in chapter 2, subjects are fundamentally minimal in the sense that they are unitarily bounded instances of formed/structured phenomenal subjectivity that has content non-relationally and intrinsically, this problem is concerned with how a phenomenality of this sort combines to realise a unified and bounded complexly structured phenomenality that has parts – that is, a non-fundamental formed subjective phenomenality which is composed of other instances of formed phenomenality. We are, at foundation, therefore attempting to move from simple phenomenal structures (i.e., instantiations of phenomenality that are spatially structured but have no parts) to a complex structured phenomenality (i.e., phenomenality constituted by a multiplicity of micro-level parts).

In what follows, I offer a means of potentially solving this problem. To achieve this, in 5.1.1 I begin by offering an outline of a theory of composition as identity that might reasonably overcome Goff's concerns regarding constitutive micropsychism outlined in section 1.3, and I employ this to guide the narrative in this section. In 5.2, I expand on this account by suggesting that an adequate solution to the first-order combination problems may be framed

in terms of ‘phenomenal dominance’. In the same section, I offer an analysis of several unsatisfactory contemporary attempts to solve the problem by employing a not wholly dissimilar notion, before offering two potentially more satisfactory means of solving the first-order combination problem(s) that might avoid the issues afflicting this aforementioned account *and* remain consistent with the stipulations outlined in 5.1.1. Finally, in 5.3, I surmise that if one of these solutions holds, we might be able to employ an account of combinatorics to explain away the palette problem, and I conclude that with these problems now solved we are left with a multiplicity of non-fundamental instantiations of unified and bounded, complexly structured subjects of experience which are potentially disposed to manifest as P-Demons *if* these structures are standing in the right reciprocal relations to their environment. In this sense, I hold that solving these first-order combination problems certainly does *not* in itself leave us with a sufficiently adequate account of how human consciousness might have occurred (because even if we are left with P-Demons, we have still not resolved the second-order combination problems or explained how these extraordinarily simplistic P-Demons might be reconciled with the complexities of human consciousness), but it does leave us with an account of relatively simplistic macro-level structured phenomenality that may very well guide us towards a solution to the second-order combination problems which, as I shall endeavour to show in chapter 6, might culminate in an account of human consciousness. By way of achieving a full account of this kind, I begin with a theory of composition as identity.

### 5.1.1 A theory of composition as identity

As I sit here wrestling with the nature of the combination problem, I am undergoing a macro-level experience of some complexity, and, if my account of constitutive micropsychism is correct, this experience is itself constituted by a complex conglomeration of micro-level

phenomenal-physical properties. If this is so, we might say that this macro-level, complex experience is constructed from micro-level, simple experiences in the same way that individual bricks come together to form complex feats of architecture. However, the problem with this account, as highlighted in my discussion of Goff's critique in section 1.3, is that the macro-level experience I am undergoing does not seem to be reducible to the micro-level experiences it is grounded in. As Goff (2017) elucidates, my experience of grappling with the combination problem does not in itself seem to be an experience held by any one of my constituents, and so Micropsychism simply cannot be true because it is not a priori clear precisely how the micro-level experiences Y may constitute my macro-level experience X. In what follows, I suggest that this problem only occurs if we endorse a truthmaking account of grounding that holds that my macro-level experience X is grounded by virtue of individual micro-level experiences Y, and in fact, the problem dissolves entirely if we endorse a theory of composition as identity which instead holds that my macro-level experience X simply *is* all X's micro-level experiential parts Y *when these parts are taken together*.

At core, Goff's critique turns on the identity commitments underpinning Constitutive Micropsychism to convey that if constitutive micropsychism is to be properly understood, it must commit itself to the reality that macro-level experiences are nothing ontologically 'over and above' (Armstrong 1997, p. 12) the micro-level experiences from which they are constituted. Yet, as Goff purports to show, our macro-level experiences are, in some sense, over and above micro-level experiences (in the sense that the micro and macro experiences are purportedly not identical), and as a result constitutive micropsychism simply cannot be true. The problem with this, as I briefly highlighted in section 1.3, is that constitutive micropsychism is not necessarily committed to an account of grounding that holds that my macro-experience X is grounded by virtue of individual micro-experiences Y. Instead, as

Chalmers (2016, p. 184) elucidates, the constitutive micropsychist might hold that the identity of my macro-experience is grounded not in the individual micro-level parts but instead in the *relationship between the micro-level parts*. In this sense, my complex macro-level experience need not be identical to just one simple micro-level constituent to the extent that the whole is identical to just one part, and instead we might suggest that the whole is identical to its parts when the parts are taken together. This would immediately diminish the force of Goff's argument because his critique is centred on the contention that, for constitutive micropsychism to hold, the identity of the micro-experiences must be one and the same as the identity of the macro-experience such that an individual instantiation of micro-phenomenality  $Y = \text{macro-phenomenality } X$ . On this account, the whole is something over and above its parts because the experiential consciousness at the macro-level is not identical to the individual instantiations of experiential consciousness at the micro-level, but on my account,  $X = Y$  would remain true if  $Y$  is taken not to denote a singular part, but all of the parts of  $X$  taken together such that the complex phenomenal whole is grounded not in singular instances of phenomenality but in the relationship between singular instances of phenomenality. Construed as such, I am advocating (in line with Baxter 1988, Armstrong 1997, Cotnoir 2013, Leibniz 1714 amongst various others) for an account of composition as identity that holds that 'mereological wholes are identical with *all their parts taken together*<sup>146</sup>' (Armstrong 1997, p. 12). Baxter (1988) elucidates this view as follows:

*'If the chair is distinct from the seat and the leg, then we are committed to co-located objects. The chair is a material object that occupies region, R. The seat and the leg are material objects that occupy region, R. This is...complete spatial overlap: there is no place that the*

---

<sup>146</sup> Emphasis added

*chair is that the seat and leg are not, and there is no place that the seat and the leg are that the chair is not. Since complete spatial co-location is unwelcome, then perhaps the seat and the leg are not distinct from the chair. (Wallace, 2011, p. 804) Suppose a man owned some land which he divides into six parcels. Over 13 come with enthusiasm for the Non-Identity view he might try to perpetrate the following scam. He sells off the six parcels while retaining ownership of the whole. That way he gets some cash while hanging on to his land. Suppose the six buyers of the parcels argue that they jointly own the whole and the original owner now owns nothing. Their argument seems right. But it suggests that the whole was not a seventh thing.’ (p. 580)*

Whilst Baxter broadly captures the spirit of this notion of identity, I specifically endorse Sider’s (2014) articulation of composition as identity in which it is not held ‘that for each part, the whole is identical to that part... rather, that the whole is identical to the parts taken together’ (p. 211), and therefore ‘the whole is nothing over and above the parts in the perfectly literal sense of being literally identical to them’ (p. 204). With this in mind, we might say that  $X = Y$ , if  $Y$  is taken to refer to the plurality of the relations between all micro-level experiences, and  $X$  is taken to refer to the combination of the relations between these micro-level experiences. In this sense, we might adopt Sider’s account by thinking of composition as identity as follows:

**Composition as Identity** If  $Y$ s combine to form  $X$ , then  $Ys = X$ <sup>147</sup>.

---

<sup>147</sup> I note that this account is slightly different to Sider’s. Sider incorporates the notion of fusion to suggest that if  $Y$ ’s fuse to produce  $X$ ,  $X = Y$ . In order to avoid confusion, I have broadened my articulation so as to frame it in terms of combination, such that  $X$  is a combination of individual instantiations of phenomenality  $Y$ . This move is deliberate and is designed to avoid a potential confusion that might arise when we attempt to reconcile terms such as fusion with one of the solutions offered to the combination problem in the next section.

In this formation the composite X and its various parts Y share an identity because, as Sider (2014) notes, “anything that is one of the Ys is part of x, and each part of x overlaps something that is one of the Ys” (p. 212).

The problem with an account of this kind, however, rests in the extent to which we might reasonably make sense of the notion that a composite is truly nothing over and over its parts. In this respect, one might hold that is not immediately, without qualification, clear that the identity of the parts taken together is equal (or the same as) the identity of the whole. This is a problem because whilst an advocate of composition as identity might reasonably hold that the whole is the same thing as the parts taken together, this does not necessarily entail that the parts and the whole share an identity, and so, as Cotnoir (2014) points out, we need an account of ‘why sameness should be considered identity’ (p. 3).

With this highlighted, as I shall elucidate in this chapter, it is perhaps possible to conceive of a very clear sense in which the macro-level instantiation of consciousness ‘Y’ *is* both numerically identical to a singular micro-level instantiation of consciousness ‘X’ *and* qualitatively identical to the relations between the various micro-experiential parts of Y without being anything ontologically over and above the relations between X and its neighbouring microexperiences, and I hold that if such an account is achievable I will have offered an account of composition that appeases both Cotnoir’s (2014) and Goff’s (2015/18) concerns. With these aims in mind, we are now left not just with the problem of providing a solution to the combination problem that explains away these concerns, but also with the problem of accounting for how certain combinations of micro-phenomenal parts seem to lead to unified conscious composites while other combinations do not. I now turn to offer a solution to the standard (and inverted) combination problem for Micropsychism.

## 5.2 Two potential solutions to the combination problem(s)

*Take a hundred of them [feelings], 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 101<sup>st</sup> feeling would be a totally new fact; the 100 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, nor (in any intelligible sense) say that they evolved it. (James 1890, p.60).*

As iterated in section 1.3, this is the standard articulation of the combination problem for panpsychism. This elucidation contains no explicit mention of subjects, but we might describe the subject-summing problem as a genus of this issue, and we might elucidate the nature of this problem by simply replacing references to ‘feelings’ in the passage above with references to ‘subjects’. On my account of phenomenal properties, we cannot separate feelings from subjects (or vice versa), and so, for my purposes, I see no reason to formally demarcate the subject-summing problem from the combination problem - I hold that, given my account of phenomenal properties, a solution to one problem will entail a solution to the other. With this established, in what follows I offer a potential solution that turns on the contention that we might make the combination problem substantially less intractable if we resist the urge to conceive of new ways in which we might make sense of a novel additional ‘101<sup>st</sup> feeling’ arising within a system composed of 100 feelings, and instead hold that there is no new feeling, there is only a dominant phenomenality (an



amplified version of one of the feelings already in the system). I propose that framing the combination problem in this way substantially reduces the difficulty involved in accounting for how a macro-level subject of experience arose amidst certain conglomerations of micro-level subjectivity, for on this account there is technically no new subject, there is only an amplified version of one of the instantiations of phenomenality that was already latent in the system, and so we immediately avoid the ontological issues involved in accounting for precisely how 100 micro-level instantiations of phenomenality produced an entirely new fact (the 101<sup>st</sup> feeling). Instead, on the view I ultimately endorse, we might say that there is at time 1 (T1) 100 micro-level instantiations of phenomenality, and, at T2, when these micro-level instances of phenomenality are related to one another in the right way, we now have 99 instantiations of micro-level phenomenality and 1 dominant instantiation of macro-level phenomenality that is an amplified version of one of the micro-level instantiations of phenomenality by virtue of, for example, remaining numerically identical to the 100<sup>th</sup> micro-level phenomenality at T1 whilst overlapping the experiential content of the 99 instantiations of micro-level phenomenality.

The problem with this, however, is accounting for precisely how this dominant phenomenality came to achieve knowledge of ‘what the other feelings are and mean’, and so framing the combination problem in terms of dominant phenomenality does not avoid the issues entailed within the combination problem. Instead, it simply changes the mode of enquiry employed to solve the problem by holding that we should attempt to solve these combinatorial issues by offering an account of how one pre-existing instantiation of formed phenomenality came to acquire a structure that would allow it to instantiate a dominant phenomenality realised as a result of the overlap of its micro-level

phenomenality with the phenomenality of other micro-level subjects. With this in mind, to clearly set the parameters of the discourse within this section, we might elucidate the nature of the combination problem (and the nature of a solution framed in terms of dominant phenomenality) pictorially:




Figure A: a 'standard' solution to combination problems

Figure B: a 'dominant' solution to combination problems

In these models, the green, blue and orange circles represent individually distinct unified and bounded instantiations of phenomenality, and the yellow/darker green circles represent a singular unified and bounded instantiation of phenomenality that has parts. Construed as such, figure A represents an entirely novel phenomenality that stands as an entirely 'new fact' that is not deducible from or reducible to the simple experiences that constitute it.

Conversely, figure B represents an amplified version of one of the instances of phenomenality that already existed in the system. In this model, this amplification occurs when one instance of phenomenality overlaps the other instances of phenomenality in the system in such a way that it perhaps achieves access to more experiential content, and simultaneously becomes more amplified (i.e., becomes richer/intensifies/grows), without introducing *entirely* novel feelings, or subjects<sup>148</sup>. In this latter model, then, there are no

<sup>148</sup> I note that one might argue that there is a sense in which if the content is now integrated into a singular phenomenal field that is experienced by a singular subject, then the content is not the same as the content that occurred prior to the onset of dominant phenomenality. I certainly sympathise with this point, but I would note that this content is still not strictly novel because we might have deduced its existence from the experiences at play in the system initially. We might say that the initial system contained experiences X, Y, Z and the

ontologically novel subjects or feelings, there are only amplified macro-level versions of the pre-existent micro-level phenomenality (and the experiential content that is already available in the system integrated into a singular whole)<sup>149</sup>. With this articulated, I note that I do not devote any time to advancing a solution that relies on the introduction of novel subjects or feelings - the literature is already overrun with solutions of this ilk (see Goff 2009/16, Seager 2010, Morch 2014 amongst others) - and I hold that the reason all of these solutions appear unsatisfying is that they all set themselves the insurmountable task of advancing a solution that, in some sense, remains consistent with 'figure A' (and therefore faces a difficulty in explaining how an entirely novel feeling/subject emerges). Contrarily, I suggest that we should attempt to advance a solution that is more consistent with 'figure B', and as such, I now turn to offer two potential ways in which we might account for the existence of dominant phenomenality.

### 5.2.1 Dominant Fusionism

The first solution I consider is framed in terms of a form of fusionism in which the dominant micro-phenomenality 'X' fuses with other instances of micro-phenomenality 'Y, Z' in such a way that the phenomenality of X is amplified to now contain the experiences of Y and Z. In this respect, the micro-phenomenality of X overlaps the micro-phenomenality of Y and Z to subsume them within it in such a way that it undergoes a phenomenal change whilst

---

combination of these is a dominant phenomenality *XYZ* if X is taken to be the dominant phenomenality that contains an amplified version of its pre-combined phenomenality that somehow now also contains the experiential content contained in Y and Z. In this sense, there is not any introduction of any novel feelings or subjects – there is merely the synthesise of the feelings and subjects that were already latent in the system. Contrarily, in figure A, we would have to say that the initial system contained X, Y, Z, and the combination of these is a novel subject with novel feelings 'G', such that G is not a priori entailed by its constituents in the same way *XYZ* is.

<sup>149</sup> We might therefore say that the dominant micro-subject (and its content) simply *is* the macro-subject by virtue of the macro-subject being simply an amplified version of the micro-subject and thereby being, in some sense, identical to the micro-subject.

remaining, in some sense, phenomenally and numerically identical to the micro-phenomenality of X that existed in the pre-fusion state of the system.

The crucial difference between this account and the form of fusionism advanced by Seager (2010) I covered in chapter 1 is that for Seager there is an entirely new subject with a novel point of view that occurs as a result of the fusion of micro-level phenomenality (and so this position falls foul to the issues involving radical or strong emergence), whereas I suggest that there is no new point of view but simply an amplified version of one of the instances of phenomenality that was already present in the system<sup>150</sup>. One relatively crude way to think of this is by employing the usage of colour (as evidenced in figure B). We might say that a pre-fusion X is a phenomenality that is simple and thereby captured by the primary colour ‘green’, and we might say that as the post-fusion X now contains the experiences Y and Z, X is now more accurately described as a different shade of green by virtue of now containing a representation of the experiences Y and Z. In this regard, the micro-phenomenality X just *is* the macro-phenomenality XYZ because on this account the micro-level ‘what it is like’ experience of X remains continuous with the macro-level experience of X whilst also now containing an amplified instantiation of X as a result of the fusion of the phenomenal fields of Y and Z. In this respect, there is no radical or strong emergence at play because there are no new subjects and also, in the strict ontological sense, no entirely novel experiential qualities or experiential content (because the content in the post-fused state remains nothing over and above a combination of the content in the pre-fused state). Instead, there is simply the fusion

---

<sup>150</sup> In this respect I use ‘fusion’ in a far weaker sense than the sense usually employed in the literature (i.e., the stronger sense in which fusion necessarily entails radical emergence), and I hold that this weaker account is consistent with constitutive micropsychism if we take it that constitutive micropsychism is only committed to the contention that micro-phenomenality somehow realises macro-phenomenality. I note, therefore, that if this is the definition employed, there is no reason why a weak account of fusion might be viewed as strictly inconsistent with constitutive micropsychism, because on the account I am offering macrophenomenality is still strictly constituted by microphenomenality (and therefore the post-fused state might simply be described as a weakly emergent phenomenon that is reducible to its constituent parts).

of Y and Z into X to necessitate a combinatory phenomenal state in which ‘what it is like’ to be Y and Z are now parts of ‘what it is like’ to be the dominant phenomenality X.

There are various ways in which we might sense of this account. We might suggest that the fusion of Y and Z into X results in a loss of the micro-level minimal subjects entailed in Y and Z whilst the phenomenal content of Y and Z is not lost. What this might mean is that the subject Y and the subject Z fuse to become a part of the dominant subject X and thereby the ‘what it is like’ experiences of Y and Z also fuse to become a part of ‘what it is like’ to be X in such a way that the pre-fused content survives the fusion whilst the pre-fused subjects do not. Alternatively, we might suggest, in line with a similar account advanced by Morch (2014), that the subjects Y and Z somehow partially survive their fusion whilst being substantially changed as a result. Such an account might lead to the suggestion that somehow the phenomenal content of Y and Z is transferred, perhaps by fusion, to X whilst the subjects Y and Z survive. The former account, then, entails the annihilation of subjects (but not necessarily their content), whereas the latter account entails the survival of subjects and the fusion of their content.

Whilst the problem of accounting for precisely how subjects of experience might cease to exist is certainly a compelling reason to query the intelligibility of the former account, I hold that the latter account faces equally difficult issues. On this account, bottom-level subjects simultaneously somehow share their content with the macro-subject whilst remaining subjects themselves. Yet, if the macro-subject is not a fusion of subjects, we seem to be left with simply a myriad of micro-subjects and, somehow, a further micro-subject with more content than its neighbours. As Nescic (2018) surmises, ‘in Morch’s account there is no strong emergence because there is the survival of microsubjects, but there is also no macrosubject in

a strong sense, and such a position is deflationary' (p. 171). Whilst I do not necessarily share Nescic's enthusiasm for an outright rejection of this idea (as we shall see, my final solution does not necessarily make use of any kind of fusion but is perhaps not entirely dissimilar in spirit), I do note that a satisfactory solution to the problem must explain how we transition from rudimentary, simple micro-phenomenality to a structured macro-phenomenality – that is a phenomenality with parts – and it is not immediately clear how we might achieve this sort of phenomenality by employing Morch's theory. More pressingly, if, as I argued in chapter 2, phenomenal content is tied to the subject, then it seems any theory akin to Morch's faces a more immediate issue, for it is not immediately clear precisely how we might make sense of the content of phenomenality Y and Z being fused into the phenomenality of X if this fusion does not also contain a fusion of subjects. In this sense, we might say that as the content of phenomenality just *is* the phenomenal subject, we cannot coherently conceive of a situation in which a subject may be separated from its content to the extent that it might survive independently from it. As such, if the phenomenal content of Y and Z is fused into X, then we would either need to also fuse the subjects Y and Z or advance an alternate account of the relationship between phenomenal content and minimal subjecthood in which we might, somehow, either conceive of a bottom-level subject of experience maintaining an existence even in the absence of experiential content *or* advance an account in which micro-subjects keep their content whilst sharing it with the dominant instantiation of phenomenality. As argued in chapter 2, I see no coherent way in which we might make sense of subjects existing independently of their content and so I devote no more time to an articulation of an account of fusionism of this kind<sup>151</sup>.

---

<sup>151</sup> Instead, I hold that one way to make sense of such an account might be to hold that minimal subjects keep their content and also, somehow, share some of their content with a dominant subject (an account of this kind shall be offered shortly)

I, therefore, argue that whilst perhaps the contention that subjects are annihilated is not overly satisfying, an account of dominant fusion that entails an annihilation of this kind offers a potential way in which we might make sense of an account of dominant phenomenality in which the core commitments underpinning a Russellian Constitutive Micropsychist account are satisfied. With this said, a dominant fusionism of this kind might leave us querying the extent to which we necessarily need to commit ourselves to the notion that subjects are somehow annihilated. We might perhaps conceive of another form of dominant phenomenality in which macro-phenomenality is achieved without relying upon this mysterious form of annihilation, and if such an account could be offered it might very well present a more explanatorily satisfying solution to the combination problem. With this in mind, I now turn to explore some of the issues associated with Kadic's (2020) and Nescic's (2018) recent attempts to advance an account of this kind, before highlighting some of the issues implicit in these solutions and using these issues as a platform from which to offer a new solution that, I argue, stands as a potentially more viable alternative to the dominant fusionism previously discussed.

### 5.2.2 Identity Panpsychism and Dominant Phenomenality

One way to overcome the unsavoury commitment to subject annihilation entailed within accounts of dominant fusionism is to posit a constitutive panpsychism predicated upon a form of identity panpsychism. As Chalmers (2017) elucidates, this sort of panpsychism is traditionally associated with Leibniz' (1714) 'dominant monad' view, in which macrosubjects are identical to a single localised 'dominant' monad:

*On this view, the subject of our experiences is a single localised fundamental entity: perhaps a single quark somewhere in our brain. The micro experiences of this quark are precisely our*

*macro experiences. There are obvious worries here about this quark's stability (what happens when it disappears?) and about its causal role (how could its properties play the rich causal role that macro experiences seem to play?).* (Chalmers 2017, p. 195)

Kadic (2020) has recently offered a form of identity panpsychism that purports to overcome Chalmers' issues. On this view, a singular, dominant, fundamental subject acts as the locus for the experiential content of other instantiations of micro-phenomenality by virtue of a group of fundamental subjects somehow routing their experiential content to this dominant fundamental subject in such a way that this micro-level singular phenomenal field acts as a conglomeration of the experiential content within a given system whilst remaining a micro-level minimal subject. In this sense, we might describe this as an account in which a multiplicity of experiential content overlaps the content of a singular fundamental subject and thereby becomes a part of 'what it is like' to be this fundamental subject. Framed as such, we might hold that this perhaps solves the combination problem and avoids the issue entailed by the annihilation of subjects of experience by holding that microsubjects survive whilst routing some of their content to a more dominant microsubject. Further, in an attempt to overcome Chalmers' issues, Kadic suggests that the dominant subject that acts as the conglomerate of the content within a given system need not be fixed, instead it might be the case that experiential content is re-routed into different fundamental subjects continuously. Whilst this might very well go some way to appeasing Chalmers, I argue that this view must do far more than simply overcome the issues highlighted by Chalmers, for all theories of this kind arguably face a more substantial problem than those previously mentioned.

The nature of this problem lies in a seeming irreconcilable disparity between our understanding of what it is to be a fundamental microphysical entity and what it means to



have, or undergo, complex experiential content, for if the fundamental subject in discussion is truly fundamental in the sense that it occupies a physical form that is absolutely simple, then it is not at all clear precisely how this simplistic form might itself act as the locus of complex experiential content, if complex experiential content is taken to refer to content that is anything more than absolutely simplistic nonsynchronous experiences.

To elucidate this, we might think of basic microphysical entities as ‘small simples’ that do not themselves have parts or intrinsic structure and cannot therefore reasonably sustain anything beyond extraordinarily simple experiences. If this view holds, there is a problem for Kadic’s account because his theory seems to require that a singular fundamental subject – a small simple – undergoes a multiplicity of experiential content simultaneously, and this is simply not something that we can easily reconcile with the concept of simple, fundamental, microphysical entities if our concept of a fundamental entity is such that this entity has no intrinsic structure and no parts. This is because Kadic is attempting to combine the experiential content ‘X’ and ‘Y’ into the fundamental microphysical experiential subject ‘Z’, but if Z’s phenomenality is just the intrinsic nature of a basic, fundamental microphysical entity it is not at all clear that Z occupies the necessary complex structure to undergo the sort of complex, synchronous experiences Kadic’s theory entails. To clarify this, we might say that as a small simple is an extraordinarily simple and basic instantiation of physicality, it can only realistically undergo one extraordinarily simple and basic experience at a time because its structure is not complex enough to sustain multiple experiences simultaneously, and so Kadic would need to account for precisely how a singular fundamental subject might undergo complex experiential content whilst remaining a fundamental microphysical entity that is devoid of parts.

It might be tempting for an advocate of Kadic's position to respond to this issue by highlighting that perhaps the fundamental subject does have parts, for its parts may very well be all of the microphysical entities that route their content to the fundamental subject. Similarly, as one proponent of a position similar to Kadic's notes, the subject does not need complex content or properties, perhaps the system as a whole just needs '[a property] for subjectivity. Subjectivity is simple. It could just need a property of being a point of view' (Nesic 2018, p. 175). The problem with this is that the subjective nature of the whole still needs to be sufficiently complex, and have sufficiently complex experiential content, to stand as a subject of the type we are trying to explain (i.e., macro-subjects of the sort associated with human consciousness). In this sense, the subject itself needs to have this content in such a way that its phenomenality is complex and rich, but if this is so we cannot rely on the notion that the subject achieves its rich, complex experiential content by virtue of the whole of which it is a part, because the parts of this whole are decidedly *not* parts of the subject itself, and so the 'point of view' of this subject must remain simple because the structure it occupies remains nothing more than a minimal fundamental instantiation of microphysicality. So, to surmise that the subject itself may have parts is to violate the commitment to identity panpsychism by inferring that the subject is now something more than a simple fundamental microphysical entity, and yet if we cannot reasonably say that the subject has parts, nor can we reasonably say that the subject has rich, experiential content.

It seems that what Kadic and Nesic are attempting to argue for, then, is a property of subjectivity that resides as a mere part of a wider system that contains other phenomenal parts, and that somehow this disunified system uses this subjectivity property to provide itself with a unified point of view that acts as a locus for the phenomenal content in the system. The problem with this, as iterated, is that for this subjectivity property to have the sort of rich,

complex phenomenal content necessary to account for our existence as macro-level subjects of experience, *this subject cannot itself be a mere simple part of a wider system but instead must itself have parts*. In this respect, we either need to offer a sufficient solution to the seemingly intractable problem of how a fundamental, simple subject devoid of parts can itself be the bed of complex phenomenal content, or we need to offer something than the form of identity panpsychism previously described by framing an account of how a singular minimal subject might come to overlap all of the phenomenal content within a given system in such a way that it transitions from a micro-level subject (a subject devoid of parts) to a macro-level subject (a subject that has parts occupies a complex structure). With this in mind, I now turn to offer a solution of this latter kind by framing a theory of ‘non-fused dominant phenomenal overlap’ which, contrary to Kadic’s view in which phenomenal content overlaps a singular dominant subject, offers an account of how a dominant phenomenal subject overlaps phenomenal content.

### 5.2.3 Non-fused Dominant Phenomenal Overlap

Therefore, contrary to Kadic (2020) and Nesic (2018) I hold that the way to solve the combination problem for constitutive micropsychism is to establish a means by which a singular instantiation of phenomenality comes to dominate other instantiations of phenomenality in such a way that this dominant phenomenality ‘Z’ comes to overlap other instantiations of phenomenality ‘Y’ and ‘X’. To elucidate this sentiment, I present two figures that represent Kadic’s view (figure A) and my own (figure B):




Figure A




Figure B

On both accounts, there is strictly no new subject that arises by way of the combination of microphenomenality, instead there is simply a pre-existent subject, or instantiation of phenomenality, that comes to contain, or dominate, the other pre-existent instantiations of phenomenality in the system. On Kadic's account, this process is achieved by way of the experiential content X and Y being communicated to the dominant subject Z in such a way that Z now contains all of the experiential content within the system whilst remaining the same experiential subject. Contrarily, on my view, the phenomenality 'Z' comes to overlap the phenomenality 'X and 'Y' and subsequently becomes a dominant, amplified instantiation of phenomenality 'Z' that 1) has parts and 2) has a representation of the content 'X' and 'Y'. In both cases, there is no fusion and thereby no need for subjects of experience to be annihilated, instead all that is required is for experiential content to, somehow, be shared with a dominant instantiation of phenomenality in such a way that this dominant phenomenality does not become a new subject but does become an amplified instantiation of one of the subjects already existent in the system. Construed as such, the crucial difference between mine and Kadic's/Nesic's views is that on my account the dominant instantiation of phenomenality 'spreads itself out' to overlap the experiential content in the system, whereas for Kadic the experiential content in the system 'spreads itself out' to overlap the dominant instantiation of phenomenality.

As iterated, this latter account faces the problem of establishing precisely how this simple dominant subject devoid of parts can have the sort of complex synchronously rich experiential content entailed by this process, whilst my solution avoids this particular problem because on this account the dominant subject overlaps its content and thereby does have parts (and thereby might rightly be construed as a macro-subject). This does not render my account devoid of potential explanatory shortfalls, however. Chiefly, it is not *prima facie* clear precisely how a singular instantiation of phenomenality comes to overlap other instantiations of phenomenality, nor is it clear how this process of overlap would lead to the experiential content of the phenomenality ‘X’ and ‘Y’ becoming a part of the phenomenality ‘Z’. In this sense, one might suggest that we must still confront the problem of accounting for how experiential qualities, or content, combines<sup>152</sup>.

To confront this problem, I suggest returning to the account of experiential content outlined in chapter 2. On this account, phenomenal content is gleaned by way of a direct self-representation of the physical structure, or form, a given instantiation of phenomenality occupies. If this is so, then we might say that if the phenomenality of ‘Z’ occupies a structure that could conceivably facilitate phenomenal overlap by either perhaps physically overlapping/copying, or somehow becoming entangled with, the physical form of phenomenality ‘Y’, then in some sense the content Y is undergoing is shared with Z because the self-representation of Z’s physicality now also contains a self-representation of Y’s form that is conceivably not entirely distinct from the content Y is undergoing. This would perhaps point to one way in which content is shared with a dominant instantiation of phenomenality, and this account holds the benefit of explaining how an instantiation of phenomenality ‘Y’

---

<sup>152</sup> I note that this particular problem is just as troubling for Kadic and Nestic.

might retain its content whilst also simultaneously sharing its content with a dominant subject 'Z' in such a way that 'Z' is amplified, or intensified, by virtue of the interaction whilst Y is not. We might think of this process of experiential intensification/amplification/overlap by way of an analogy:

*'A light wave can hit other atoms and force an increase in the light intensity. By such processes, the light waves reach certain amplitudes. Haken (1987, p. 128) says that one light wave 'enslaves' the others; this means that it becomes dominant and orders the system. As a result, an ordered light wave – the laser beam – emerges. From a chaos of light waves, an ordered pattern emerges...A light wave is caused by a fluctuation – that is, an electron returns to the inner trajectory and emits energy; a light wave can intensify itself by 'enslaving' electrons. Such an intensification always means circular causality, because one entity causes the behaviour of another entity, and this behaviour results in a transformation of the first entity. Due to such intensification, the system enters a system of chaos/instability/bifurcation. A certain light wave is selected and determines the emergence of the laser beam. It is determined that a laser beam will emerge, that fluctuations and intensification will result; but not determined is exactly how this will take place and which light wave will order the system'. (Fuchs 2003, p. 7)*

In this respect, as with the dominant light wave, as one instantiation of phenomenality comes to dominate the others, it is amplified and intensified by virtue of these interactions, whilst conceivably the instantiations of phenomenality it interacts with are lessened somewhat as they become 'enslaved' or 'overlapped' by the dominant phenomenality of which they are now a part. Whilst this analogy might go some way to elucidating the core of this idea, and whilst we might concede that perhaps experiential content could be shared in the manner I

describe, the account I am offering remains explanatorily suspect without a sufficient theory of precisely how a simple microphenomenality might ever have come to overlap other instantiations of simple microphenomenality in the first place. With this in mind, I now turn to offer a theory of this kind.

One way to make sense of this is by extrapolating a new take on Goff's (2016) 'phenomenal bonding relation'<sup>153</sup> in which we conceptualise it not as the relation that 'when two subjects stand in it, they produce a further subject' (p. 293), but instead as the relation that 'when two instances of phenomenality stand in it a dominant and submissive phenomenality will occur' (this is what I call a 'dominant phenomenal bonding relation'). I propose that one way of making sense of this might be advanced by thinking of the absorption and expulsion of exchange particles as the bonds that mediate the sort of relation I am discussing. At the bottom-level, we might think of gluons as mediums by which a dominance of this kind occurs because they connect quarks together (although I hold that all exchange particles might very well act as mediums for a bonding relation of this kind). In this sense, I suggest that as an exchange particle connects one entity to another, the phenomenality of one entity either overlaps the phenomenality of the other entity (and thereby dominates it), or it does not. I am reticent to make any explicit claims about precisely how this process might occur, but I tentatively note that perhaps when standing in the right relation, dominance might either occur as a result of the opposite sort of process associated with diffusion, such that instead of

---

<sup>153</sup> Chalmers (2016) described this view as follows: 'microsubjects constitute macrosubjects in virtue of certain phenomenal relations between the microsubjects: phenomenal bonding relations', and holds, as I do, that because 1) physics captures mere mathematical structure and not the intrinsic nature of this structure and 2) microphysics postulates spatial-temporal relations, then, when contextualised in terms of Russellian Micropsychism, 'it is not out of the question that a certain phenomenal relation could serve as the quiddity underlying spatiotemporal relations' (p. 23-24).

two gases mixing to spread out evenly over an available space, we might say that when two instances of phenomenality meet, the phenomenality with the most mass comes to dominate the phenomenality with the least mass to spread out and occupy more of the available space as the submissive phenomenality is overlapped to the extent that it shrinks without being annihilated<sup>154</sup>. Whilst this sort of phenomenal interaction might be in principle conceivable by introducing an appropriate conception of exchange particles acting as the medium with which otherwise isolated instantiations of phenomenality might interact by virtue of being physically conjoined, one potential problem with this particular view might be that we arguably reintroduce the mystery of precisely how phenomenality might interact with other instances of phenomenality if both phenomenal fields remain properly ‘unified and bounded’. Conversely, then, we might say that there is no strict interaction between instances of phenomenality and instead one instance ‘Z’ simply comes to physically overlap another instance ‘Y’ to achieve a representation of the content of Y by way of overlapping the physical form that Y occupies. On this latter account, there is no phenomenal interaction at all but there is still phenomenal amplification as the phenomenality of Z grows as it overlaps, exchanges particles, and continues the process of overlap. In either of these cases, I suggest that we can conceive of how these processes might leave us with a unified whole that instantiates a bounded, amplified version of a dominant phenomenality that overlaps, or dominates, other instantiations of phenomenality in such a way that it now has parts and complex content. Further, I hold that we might conceive of this because all this process entails is the commitment to the potential that phenomenality might come to spread itself and

---

<sup>154</sup> We might think of this in terms of, for example, the interaction between a top quark (a quark with an estimated 91,000 mega electron volts) and a bottom quark (a quark with an estimated 4,800 mega electron volts). We might hypothesise that the dominant phenomenal bonding relation between these two quarks might result in the phenomenality associated with the top quark dominating the phenomenality associated with the bottom quark, because the mass of the former is greater than the latter. I note, however, that this is simply one possible way of conceiving of the dominant phenomenal bonding relation that is itself highly theoretical.



overlap other instances of phenomenality in such a way that it remains numerically identical to the phenomenality 'Z' that was available in the pre-overlapped state whilst now existing as an amplified version of 'Z' (what we might term  $Z_{xy}$  if 'x' and 'y' are taken to refer to parts of Z/experiential content that Z has overlapped), much in the same way that the laser beam might be thought of as an amplified, or intensified, version of one dominant light wave after this dominant light wave has acquired various parts<sup>155</sup>. Further still, much in the same way that the causal powers associated with the laser beam increase in scope as the beam increases in intensity, we might say that so too do the causal powers associated with phenomenality increase in scope as the phenomenality acquires new parts and becomes increasingly amplified<sup>156</sup>.

---

<sup>155</sup> I note that, in this model, the dominant instantiation of phenomenality might very well persist in time by constantly acquiring new parts. With this in mind, it is of interest to note here the recent confirmation of the hypothesis (see Bergami & Berninger 2012 and Moreno & Rhiner 2014) that as the brain forms new neurons, these neurons are selected on a competitive basis as they attempt to co-exist with their older counterparts and form synapses with the pre-existing network already in place until these newer neurons either fully mature and achieve integration into the wider network as their older (or less optimal) counterparts die, or are themselves eliminated (p. See Bergami & Berninger 2012, p. 1016). In this case, this might be a physical representation of the process of the dominant instantiation of phenomenality acquiring a new part and overlapping it to the extent that is sufficient in order to maintain its form and existence by continually acquiring new parts (neurons) that are selected for on the basis of fitness. It is of equal interest to note the recent finding that suggests there might be a giant neuron that wraps around the entirety of a mouse brain (see Reardon 2017). Koch (2005/17) has suggested that this same phenomenon might be detected in humans, and that this neuron might extend from the claustrum – an area Koch maintains might act as the locus of consciousness in mammals. If these findings are replicated, they might very well serve to substantiate my appeal to dominant instantiations of phenomenality coming to acquire new content by overlapping their counterparts, for just as 'from the outside' the claustrum might send out 'neuronal tendrils' so as to acquire access to new information pathways in the brain, 'from the inside' the claustrum might itself be a dominant instantiation of phenomenality that is spreading itself over the rest of the brain in a bid to acquire the parts necessary in order to integrate experiences into an informationally rich unified whole that might more effectively navigate its environment and drive behaviour.

<sup>156</sup> In this respect, we might have appeased the concerns of Goff (2017) (see chapter 1), for we now have an account of composition that explains how this amplified instantiation of phenomenality might be numerically identical to a certain micro-level instantiation whilst also acquiring various parts that instil various additional causal powers and qualities without manifesting anything ontologically over and above a collection of the qualities/powers entailed by microphysics and microexperience, or the qualities/powers that supervene on the interactions between microexperiences and microphysics. In this respect, we have arguably addressed what Goff (2017) terms 'the strange middle way between identity and distinctness...[In which] fact X involve[s] different objects and properties to fact Y, and yet add[s] nothing beyond the objects of properties already involved in Y'(p. 382), because, on this account, fact X is both qualitatively distinct and numerically identical to one of the instantiations of Y, by virtue of being qualitatively identical to all its constituents taken together in such a way that X is something over and above the properties entailed by this singular instantiation of Y, but is not something over and above this singular instantiation when this instantiation overlaps its neighbours. Therefore, if 'Y' in this account is taken to refer to microexperiences and the relations that hold between them (i.e., microphysics), then it is not difficult to conceive of how  $X = Y$ .

With all of this said, I note that a solution to the ‘inverted’ combination problem should not pose any issue given that we might simply suggest that those systems that do not seem to exhibit a singular unified and bounded experiential whole are simply those systems that did not stand in the sufficient dominant bonding relation to manifest the sort of phenomenal dominance that might lead to a unified and bounded experiential whole<sup>157</sup>.

I therefore end this section by reiterating that upon curtailing our potential solutions to the combination problem to the extent that we only concern ourselves with solutions that do *not* entail the possibility of ontologically novel instantiations of phenomenality arising as a result of the combination of micro-level phenomenality, we begin a line of enquiry that avoids many of the issues typically associated with the first-order combination problem(s). In light of this, I note that whilst dominant fusionism or non-fused dominant phenomenal overlap are certainly not devoid of certain explanatory shortfalls and ambiguities, these ambiguities might very well occur as a result of a broader issue afflicting all theories of phenomenal combination, which occur because all face the problems entailed in attempting to map how phenomenal combination might occur in a way that is broad enough to be made consistent with the accounts of combination advanced in the physical sciences. As a result, the majority of purported solutions make use of broad analogies in an attempt to satisfactorily elucidate a potential solution that might remain broadly consistent with our understanding of the natural world without unnecessarily overextending themselves by offering explicit details of the precise mechanisms involved in the formation of phenomenally unified composites. At least in this respect, my formation is no different from other ostensive solutions in so far as it

---

<sup>157</sup> Conversely, if the former outline of dominant fusionism is endorsed, we might simply suggest that the issue of non-experiential wholes is explained away as the consequence of those composites that did not undergo phenomenal fusion of the necessary kind to manifest as unified and bounded wholes.

offers a broad theory of how phenomenal combination might occur without elucidating specific details.

Whilst this is certainly a potential explanatory shortfall that one might deem almost inevitable given the theoretical nature of any potential solution to the combination problem, I do note that the formation I am offering does remain broadly consistent with our understanding of how physical composition occurs. This is the case because whilst we might traditionally think of composites forming as a result of atoms being combined to create new molecules, we might also suggest that perhaps ‘molecules do not need to fully blend into each other to create a macromolecule; a cell is a grouping of macromolecules but when we move from lower level micro-states to higher level microstates the lower level micro-states retain their individuality to a degree’, and in this sense ‘there is not total overlap and each retains its identity to a degree in the bond; yet they do marry’ (Montero 2016, p. 224). As Montero suggests, if we conceive of composites in this way, there is no good reason why the same should not be true for phenomenality, because ‘for the panpsychist, micro minds exist as distinct entities at the microlevel, but at higher levels they undergo some changes and, while retaining their identity to a degree, enter into the everlasting bond of marriage’ (p. 224)<sup>158</sup>. Construed as such, the appeal of dominant phenomenality is rendered readily apparent because the ‘marriage’ Montero speaks of becomes simply another way of thinking of the dominant bonding relation (or the dominant fusion) that occurs as one instantiation of phenomenality becomes everlastingly amplified as a result of its interactions with its neighbours. If this is the case, I hold that whilst neither of the solutions I have presented is

---

<sup>158</sup> In this respect the distinction between dominant fusion and non-fused dominant phenomenal overlap might be conceived of in terms of the type of identity that persists. In the former case, there is a sense in which the identity of one pre-fusion subject and all of the pre-fusion experiential content seems to persist in the post-fused state, in the latter case there is a sense in which the identities of all of the pre-overlap experiential content *and* subjects persist in the post-overlap state.

beyond reproach, it does appear at least conceivable that either one of the accounts of phenomenal dominance articulated in this section might at least act as a reasonable candidate for a solution to the first-order combination problems<sup>159</sup>. Having established this, I now turn to take stock of what we might be left with, and what other solutions may arise, if either one of these solutions is adopted.

### **5.3 A consequence of these solutions: A response to the palette problem**

According to contemporary physics, at the bottom-level there is a multitude of elementary particles which can be disambiguated into a relatively limited set of fundamental physical properties. When we combine this account with micropsychism, we might say that there are as many fundamental unified and bounded subjects as there are fundamental particulars and hold that the qualities (or content) associated with these subjects are relatively limited at the bottom-level if the physical properties at the bottom-level are themselves limited. We might even go so far as to suggest that as there are, according to contemporary physics, only 12 fundamental types of particles, the universe at foundation contains a rather limited ‘palette’ of experiential content because, as argued in chapter 2, if experiential content is a direct self-representation of the physical form the subject takes, then as the fundamental forms of the physical are limited at the bottom-level so too is the experiential content. This is one way to articulate the ‘palette’ problem for panpsychism, which essentially amounts to the issue of accounting for how we achieve the sort of rich, complex, varied experiential content we undergo if the universe is fundamentally phenomenally sparse.

---

<sup>159</sup> Although I note that non-fused dominant phenomenal overlap avoids the annihilation of subjects and thereby prima facie appears preferable to dominant fusionism.

I hold that with the problem articulated as such, either one of the solutions to the first-order combination problem(s) I have offered would stand as a suitable candidate to resolve this particular issue, and I suggest that we might achieve this by simply combining the theory of dominant phenomenality with a theory of combinatorics. On this view, we might say that as the dominant instantiation of phenomenality comes to either overlap or fuse, the experiential content it undergoes is a result of the particular manner, or arrangement, in which the fusion or overlap occurred. In this sense, it seems at least conceivable that there are many ways in which the bottom-level experiential content might be arranged, and each of these specific arrangements might be associated with a distinct type of amplified phenomenality when they are appropriately integrated as parts of a dominant subject. If this is so then if we start with 12 distinct instantiations of experiential content, and we hold that each particular arrangement of this experiential content gives rise to a distinct form of amplified phenomenality when it is overlapped by or fused into a dominant subject, then it seems at least possible to apply basic combinatorics to this process and hold that much in the same way that if no repetition is granted there are 479001600 permutations of 12 digits, there are conceivably an equal number of ways to combine 12 instantiations of bottom-level experiential content. Further, given that each combination might reasonably yield a distinct form of amplified phenomenality that is differentiated from other such forms, we might even extrapolate that the process of combining, or arranging, these fundamentally sparse phenomenal qualities could potentially give rise to an infinite set in which each combination leads to a potential further combination, and thereby provides a basis from which to explain increasingly rich, varied and complex phenomenal content.

If this is the case, I suggest that upon employing an account of dominant phenomenality in an attempt to resolve the first-order combination problems, we might reasonably be left with

both a wide variety of distinct instantiations of amplified, structured phenomenality and also, as I shall now argue, with certain structures in the necessary conditions to satisfy the manifestation of, what I termed in chapter 4, a ‘P-Demon’. With this in mind, I suggest that whilst the solution to the combination problem I have offered might very well go some way to explaining why it is that we exist as macro-level unified and bounded subjects of experience, I hold that prior to the onset of the sort of macro-level experience associated with human subjects, this sort of dominant phenomenality is likely to have occurred in less complex physical, macro-level structures. And, if this so, it seems conceivable that some of these structures, by virtue of existing as a singular dominant subject of experience that has various parts, would now (in at least some iterations) satisfy the conditions to behave like a P-Demon<sup>160</sup>. This is because, as iterated in chapter 4, these structured subjects are disposed to individuate between the experiential content they undergo in such a way that this phenomenality might itself be employed to sustain a given structure in a non-equilibrium state, and so if 1) this is possible and 2) dominant phenomenality is embraced a solution to the combination problem, it is conceivable that we might be left with at least some instantiations of structured phenomenality which, when occupying the right kind of structure and standing in the right reciprocal relation to its environment, might reasonably be said to be disposed to employ their phenomenality to sustain themselves in a state of thermodynamic disequilibrium and thereby manifest as what I have termed a P-Demon. With this in mind, I now turn to the penultimate chapter of this thesis, in which I attempt to take this concept of basic P-Demons and reconcile it with Bejan’s (1997-2016) Constructal Law (and an account of thermodynamic natural selection) to display how the evolution of the disposition to phenomenally individuate might lead us to a satisfactory solution to the structural mismatch

---

<sup>160</sup> See p. 225 for an account of these conditions.

problem, the awareness problem and, most importantly, bring us one step closer to a 'complete micropsychist theory of consciousness'.

## Chapter 6

### Towards a Complete Micropsychist Theory of Consciousness

Bejan's (1997 - 2016) Constructal Law makes a clear claim about the time-direction of the universe. It suggests with no uncertainty that for a system to survive (to persist in time) it must arrange itself as to access increasingly energetic/structural states (or, when framed intrinsically and in terms of the account of Russellian micropsychism I have offered, increasingly experiential states). If this is the case, and it is at least conceivable that in certain structures this behaviour is grounded in the power of phenomenal individuation that manifests because of the quality of structured phenomenality, then as soon as 1) phenomenality becomes complexly structured (i.e., has parts), and 2) the reciprocal relations at play in the universe facilitate the manifestation of this behaviour, we might describe the evolution of this behaviour in terms of thermodynamic natural selection.

Darwin's theory of evolution as natural selection turns on the notion that as biological systems compete for limited resources, the most effective, or 'fittest', mutations within these systems will be those that are selected for survival (Darwin & Wallace 1858). Recently, certain theoretical physicists (see Kaila & Annala 2008, Smith 2008, Chudnovsky 1985) have attempted to reconcile this evolutionary principle with the principles underpinning thermodynamics. On some of these accounts (see specifically Kaila & Annala 2008) it is held that there is no obvious disparity between the fitness criterion, i.e., select the 'fittest' unit and the principle of least action, i.e., take the shortest path, for both principles may find reconciliation when framed in terms of the second law of thermodynamics<sup>161</sup>.

---

<sup>161</sup> Because selecting on the basis of the 'fittest' unit or shortest path both maximise the production of entropy when framed in terms of the system as a whole.



For my purposes, I do not deem it necessary to elucidate precisely why these two principles may be unified within the second law. Instead, I take it to be a relatively uncontroversial claim that any living system that takes the ‘shortest’ path towards more energetic states will be selected as the ‘fittest’ system, on the basis that the shortest path always results in greater access to more energetic states and always results in less energy expended in acquiring this state. As this process continues, it is not difficult to see how the population of those systems which most effectively and consistently take the ‘shortest’ path will increase at the expense of their neighbours by virtue of navigating their environment more effectively and hoarding the finite resources (i.e., energy) necessary for them to continue to persist in time. Kailia & Annala (2008) explain this in terms of biological evolution as follows:

*‘In a biological context, when two rather similar species (i.e., energy transduction mechanisms) compete for the same source of energy (e.g., food), the one with even slightly more effective mechanisms (e.g., claws, teeth, feet, etc.) captures more than the other. Gradually, the population of the more effective species will increase at the expense of the other. The overall process is pictured as flows of energy that gradually and naturally select the more effective, steeper paths. In biology, this physical consequence, which can be deduced from Lyapunov stability criterion, is known as the competitive exclusion principle.’*

(p. 2)

With this said, in line with my account of P-Demons offered in chapter 4, I hold that it is not unreasonable to think of phenomenal individuation as a mechanism, not wholly unlike claws, or teeth, or feet, that is conceivably employed to ‘capture’ more energy, and subsequently suggest that this basic informational mechanism should be subject to the impositions of

thermodynamic natural selection by virtue of only the most effective mutation of the disposition to phenomenally individuate achieving more access to increasingly energetic/structural/experiential states and thereby persisting in time.

To succinctly capture this sentiment, I employ the following maxim (derived from Bejan's 1997 - 2016 work) to guide the discourse that follows:

*The extent of a given system's survival is dependent on the extent to which that system can structure itself to provide greater access (i.e., the 'shortest path') to the energy flows that run through it<sup>162</sup>.*

With this as grounding, and if we hold that 1) if the reciprocal relations at play allow, structured phenomenal subjectivity might very well manifest like a P-Demon by employing the disposition to phenomenally individuate to guide the structures that it occupies towards more energetic states, and 2) the dispositions of P-Demons are integrated into the causally closed physical world and at least some of the pressures acting upon the causally closed physical world are evolutionary pressures, then we might expand on the above maxim as follows:

---

<sup>162</sup> This is a rephrasing of Bejan's Constructal Law: 'for a finite flow system to persist in time (to live) it must evolve such that it provides greater and greater access to the currents that run through it' (Bejan 2010, p. 1335). I take this rephrasing to amount to an uncontroversial claim given that it is equally uncontroversial to claim that all living systems are thermodynamic non-equilibrium systems in which there are sustained net flows of matter/energy either within the system or between the system and its environment (see Prigogine 1993, Popovic 2017, Barge et al 2017 amongst various others), and therefore a system at thermodynamic equilibrium (or a 'dead' system) is a system in which there are no net macroscopic flows of energy or matter either within a system or between a system and its environment. Subsequently, for a system to avoid thermodynamic equilibrium and thereby 'persist in time', it must arrange itself to provide more access (i.e., more access than its neighbours) to the energy flows that run through it.

*The extent of a given system's survival is dependent on the extent to which that system is capable of structuring itself as to provide greater access to the energy flows that run through it, and subsequently, in the case of those systems disposed to act like P-Demons, if the disposition to phenomenally individuate is found to confer a survival advantage<sup>163</sup>, a system's continued survival will be dependent, at least in part, upon the evolution of the system's capacity for accurate and efficient phenomenal individuation between the structural states that represent these energy flows.*

This is conceivably the case because as soon as P-Demons are integrated into our naturalised account of physical causation, and we admit evolutionarily pressures into this account, there is a 'competition' between P-Demons as they are conceivably selected against based on, at least in part, the extent to which a given instantiation can employ accurate and efficient phenomenal individuation to restructure itself in the direction of more energy flows/experiences/structural changes. I hold that if the P-Demon is *not* constrained by evolutionary pressures of this type, one might reasonably push for an account of precisely why this would be the case given that, on the naturalised account I have offered, P-Demons are clearly integrated into the causally closed physical world in such a way that, if these pressures exist<sup>164</sup>, P-Demons are *not* beyond their purview. I hold that *if* this is denied, the denial must either turn on an attempt to discredit the theoretical virtues of Meta-Darwinian survival of the fittest or (somehow) turn on a rejection of the claim that phenomenal

---

<sup>163</sup> I hold that that given that equipping a system with this disposition imbues at least some informational content that might reasonably be employed to situate the P-Demon in the world and direct itself towards more energy flows, it appears likely that those systems not equipped with this disposition appear to be at an informational disadvantage that is likely to confer a survival disadvantage.

<sup>164</sup> As a result of my commitment to explanatory power as a tool with which to assess the virtue of a given naturalised theory, I hold that, in absence to evidence to the contrary, these pressures *do* exist because admitting them into our account of reality provides a distinct explanatory advantage that is not shared with those accounts that deny the reality of these pressures (i.e., it provides an account that explains the development of life and is consistent with the available palaeontological data).

individuation is now a part of the causally closed physical world. I take it the latter is difficult to hold given the account of Russellian Micropsychism I have offered, and I take it that the former is suspect when contextualised in terms of both the commitments to naturalism outlined in chapter 3 and the explanatory power of natural selection as a theory about one of the pressures acting upon certain systems in the causally closed physical world that is both consistent with paleontological records and capable of explaining the development of life on earth<sup>165</sup>. Therefore, if the time-direction of the P-Demon *is* curtailed by such pressures, those P-Demons that survive will be those that have evolved a mode of phenomenal individuation that is optimised to the environments in which they find themselves, and this is the case because phenomenal individuation conceivably plays an integral role in restructuring certain systems to ensure the continued access to energy flows necessary to both maintain structured stability and adjust to conditions that are optimal to survival (i.e., achieve homeostasis)<sup>166</sup>.

In what follows, I build upon the worldview I constructed in the previous chapters to advance an account of the evolution of P-Demons. I argue that such an account may very well be employed to solve the second-order combination problems, and expand the tools we might employ to explain away the ‘inverted’ combination problem, because the evolution of P-Demons precipitates structured phenomenality that manifests either within or between two disparate states of physical systems: unchanging ‘dead’ systems that are minimally experiential at a micro-level by way of only undergoing their self-representation of their own unchanging micro-structures, and changing ‘alive’ systems that are maximally experiential at

---

<sup>165</sup> See Torday & Rehan 2017 for a broad survey of the strengths of this theory and see Fowler 2007 for an account of the theoretical virtues of naturalised Meta-Darwinianism.

<sup>166</sup> Structured stability here refers to a state of homeostasis, in which constant changes occur but some uniformity to the system or essential structure is maintained (Torday 2015 refers to this as the ‘mechanism of evolution’ (p. 573), such that evolution is reduced to the evolution of homeostasis). I take it that some available free energy is necessary to consistently sustain this homeostasis and thereby survive, and therefore the mechanism that is employed to access, or select between, energy flows will be subject to the impositions of evolution.

a macro-level by way of sustainably undergoing a unified phenomenal self-representation of a changing macro-level structure. In this sense, I maintain that the evolution of effective P-Demons results in structured phenomenality manifesting in various intensities throughout space-time.

With this as grounding, I argue that we might, in principle, employ the concept of P-Demons to explain our own conscious states but hold that, in order to explain the manifestation of the structured phenomenality associated with human subjects, we must first elucidate a potential series of evolutionary ‘jumps’ that the Demon grounded in structured phenomenality undergoes. Further, I maintain that once we have offered a general principle that might guide these various ‘jumps’ (a principle of the type elucidated above and reiterated below) we will have perhaps provided a means with which to solve the second-order combination problems. And, further still, I suggest that in so doing might even lay the foundation from which to advance a theory of, what I term, a ‘phenomenality of thought’ that might typify both an inevitable evolutionary consequence of P-Demons and lay the foundation from which to offer a complete micropsychist theory of consciousness capable of providing some insight into one potential way with which we might conceive of the relationship between phenomenality and cognition. Given this, I now turn to address the second-order combination problems (and lay the foundation from which to erect this ‘phenomenality of thought’) by addressing the intricacies and ultimate consequences of the maxim: *for a P-Demon to persist in time it must evolve such that it optimises its capacity for phenomenal individuation.*

## 6.1 The evolution of P-Demons: A solution to the second-order combination problems

*The P-Demon (the self-sustaining dominant phenomenality) is the silent pianist sometimes playing, sometimes hearing, the piano of life as the environment vibrates at various frequencies and instils various experiences in the pianist as a result. The P-Demon, as with the pianist, over time learns to differentiate between these various experiences and thereby learns to play the piano in such a way to perpetuate its access to life's music.*

If the maxim holds, we might suggest that the evolution of effective P-Demons runs such that, as a P-Demon persists in time, its capacity to phenomenally individuate between phenomenal content increases in accuracy, complexity, and efficiency so that it might individuate between content in a manner that facilitates the 'shortest path' towards more energetic states.

If this is the case, amidst this process of optimising the disposition to phenomenally individuate there seems to be, at least in some sense, optimisation of a dominant phenomenal subject's capacity to delineate the disparities, or lack thereof, between phenomenal content. Subsequently, we might say that there is a reduction to the amount of, what Friston (2005/09/12) might term, phenomenal 'surprises', or what Shannon (1948) might term 'informational entropy', relative to the content that the dominant subject undergoes. In what follows, I elucidate the intricacies of what might occur if a P-Demon evolves to reduce the amount of informational entropy, or the amount of 'phenomenal surprises/uncertainty', relative to its content, before concluding that if P-Demons did evolve in this manner, this

might very well go some way to solving the second-order combination problems and provide a grounding from which to advance a theory of phenomenality of thought that might be exercised to reveal a new take on cognitive phenomenology (and subsequently a new take on the nature of human cognition).

### 6.1.1 Informational Entropy and Phenomenal Surprise

As you occupy this moment in space-time you are observing a myriad of symbols on a page, and you can decipher the meaning contained within the disparities between these symbols because the informational entropy (or the level of uncertainty) relative to the stimuli you are undergoing is relatively low. This inference might be explained by way of Shannon's (1948) contention that the entropy of a given stimulus can be used as a measure for information one might extract from said stimulus. What we term 'information entropy' in this context is therefore a measure of the possible informational states a given stimulus might occupy, such that if the stimulus is highly 'informationally entropic' the amount of information we have relative to this stimulus is low. Conversely, if the stimulus is of low informational entropy relative to us, the amount of information we have about the stimulus is high. In this sense, this concept is not dissimilar to the amount of perceived 'surprise', or uncertainty, we might undergo relative to a given stimulus, in that sense that if the informational entropy (or surprise) relative to some content is high, the information we hold about said content is low (and vice versa).

To elucidate the notion of information entropy with more clarity, let us consider flipping a coin. For our purposes, let us imagine a human being of average intelligence<sup>167</sup> 'S' that had at

---

<sup>167</sup> In this context, intelligence entails a capacity for cognition that 'involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience' (Gottfredson 1997, p. 13).

time 1 (T1) never seen a coin, and then imagine performing (at T2) a fair coin toss in their presence. From the perspective of ‘S’, the informational entropy relative to the coin would be greater at T1 than at T2, and this would be because, in Shannon’s model, the amount of information that S might extract from the coin diminishes in time as the variables (i.e., the possible states the coin might occupy) inherent in the stimulus are exhausted. In this sense, the informational entropy of a given variable might be described in terms of unpredictability, uncertainty or ‘surprise’, and I propose that this concept might be employed to strengthen our understanding of P-Demons because much in the same way that, in line with the reduction of the informational entropy relative to the coin, the subject ‘S’ can individuate between the content it undergoes more effectively and accurately at T2 than at T1, the P-Demon’s capacity for phenomenal individuation might so too become optimised as, over time, the P-Demon evolves structures that facilitate the process of reducing the informational entropy relative to its phenomenal content.

From this perspective, we might ascribe this same concept to P-Demons and phenomenal content by thinking of this in terms of what I term ‘phenomenal surprise’ – that is, a phenomenal state that is informationally entropic relative to the P-Demon<sup>168</sup>. Further, in line with Schwartenbeck et al (2013) and Friston et al (2009), we might hold that this notion of surprise is essential to our understanding of how systems (or, as I argue, by extension P-Demons) might persist in time, because ‘in order to underwrite its existence and avoid the dispersion of its physical states...[a system] has to minimise its long-term average surprise (i.e. Shannon entropy) over the states it visits. Surprise is an information-theoretic quantity...Every system that maintains itself conforms to the imperative of minimizing the

---

<sup>168</sup> | note that surprise in this context merely denotes a phenomenal state that is novel relative to the subject. In this sense, I employ the term relatively loosely to denote a phenomenal state that, as I elucidate in the subsequent section, is not sufficiently ‘amplified’.



surprise associated with the states it encounters' (Schwartenbeck et al 2013, p. 1). Subsequently, given that an increase in the number of phenomenal surprises, or the informational entropy relative to phenomenal states, is likely to result in a decrease in the expediency and accuracy with which a P-Demon might individuate between its phenomenal states and thereby access the 'shortest' path towards more energetic states, we might say that as a P-Demon undergoes its phenomenal content it evolves to form a structure that will facilitate a reduction in the amount of 'surprise', or informational entropy, that accompanies the phenomenal content it undergoes. Put simply: A P-Demon that can delineate the disparities, or similarities, between phenomenal content with more efficacy is more likely to persist in time, and this process is only achieved by reducing the informational entropy relative to its phenomenal content.

With this contention in mind, the task is now to delineate precisely how we might transition from simplistic P-Demons with a basic capacity for phenomenal individuation to complex subjects who hold the capacity to undergo complex phenomenal content, recall similarities/disparities in variables and reduce the informational entropy residing in this content with rapidity. To achieve this, I now turn to Gabora's (2002) and Kauffman's (1993) work to offer an outline of one potential way in which P-Demons might have evolved to eventually become subjects of the type we identify with.

### 6.1.2 Amplifying Phenomenal Information

Gabora (2002) speculates that one way to make sense of simplistic instantiations of phenomenality becoming complex conscious subjects is to make use of the concept of amplified phenomenal information. On this account, the inference that information has a phenomenal component is taken as a postulate and from there it is surmised that as a physical

system amplifies its informational content it also amplifies its phenomenal content, such that the greater the degree to which information is locally amplified, the greater the degree to which consciousness is present (Gabora 2002, p. 3-4)<sup>169</sup>. Whilst I am sympathetic to Gabora's overarching aims, I note that for my purposes we only need to delineate one potential means by which a dominant P-Demon might reduce the informational entropy relative to its phenomenal content, and therefore I need not necessarily commit myself to a theory that might lead to a dual-aspect theory of information of the type adopted by Gabora (2002) or Chalmers (1996)<sup>170</sup>.

Instead, I propose adapting the sentiments espoused by Gabora (2002) in order to contend that the greater the extent to which the informational entropy relative to a P-Demon's phenomenal content is reduced, the more amplified that P-Demon's consciousness will be *and* the more efficacious its capacity for phenomenal individuation will become. I substantiate this on the grounds that in order for the informational entropy relative to phenomenal content to be reduced, the P-Demon must evolve such that it occupies a complex structure capable of facilitating the amplification of phenomenal content in a way that also amplifies the information conveyed by this content, and therefore as its structure increases so too does its consciousness *and*, if the structure is optimised to reduce informational entropy (i.e., optimised to maximise the amount of information conveyed by some phenomenal content), so too does its capacity for phenomenal individuation. To outline how a structure of


---

<sup>169</sup> In this sense I note a similarity to Tononi (2004), who posits that that the extent of the consciousness in a system might be relative to the extent of integrated information.

<sup>170</sup> Although I note that there is a very clear sense in which my formulation entails that if information has a physical component it also has a phenomenal component, I suggest that whilst Chalmers's dual aspect theory of information might lead us to the inference that information is itself fundamental (see Chalmers 1996, p. 287), my formulation of Russellian micropsychism entails no such commitment. I take this to be a core disparity between those, like Gabora (2002), Tononi (2004) and Chalmers (1996), who appear in places to identify consciousness with information to the extent that if information is not present consciousness cannot be, and myself, who identifies consciousness merely with certain instantiations of physicality (specifically, all fundamental instantiations and certain higher order instantiations).

this kind might occur, and how such a structure might amplify phenomenal content in the manner necessary to reduce phenomenal surprise, I propose, in line with Gabora (2002), turning to Kauffman's (1993) work on autocatalytic structures.

Kauffman's work was offered as a solution to the problem of life, but, as Gabora (2002) points out, it might just as easily be applied to consciousness. On this account, we might consider 'self-organising catalytic polymers [as] an information trapping web' (Gabora 2002, p. 6), and from there we might suggest that autocatalytic networks such as these hold the key to reducing phenomenal surprise by amplifying phenomenal information. To elucidate this idea, Gabora (2002) employs the following figure:



*'Figure 3. (a) A set of loose buttons. (b) Tie two randomly chosen buttons together with a thread. (c) Repeat over and over. Occasionally lift a button and see how many connected buttons get lifted. (d) Increasingly large clusters emerge, and eventually reach a point where they form one giant cluster containing most of the buttons.'* (p. 7).

For our purposes, an intricate account of the nature of autocatalytic structures is of little import, but what is of import is the basic concept captured above that, over time, connections between disparate structures might form to create a web that binds these structures into a unified whole in such a way that once one connected structure is accessed, the entire web is also accessed (for our purposes autocatalytic structures might simply stand as one observable

example of this process). With this in mind, we might conceive of a system of autocatalytic structures that are accessed by the dominant P-Demon in order to more effectively navigate its environment. In this model, we might think of the dominant P-Demon as overlapping each of its constituents until it forms a web-like structure similar in kind to the one described, and then imagine a myriad of smaller web-like structures branching from each of its constituents until we conceive of a dominant autocatalytic structure that contains a myriad of smaller autocatalytic structures that are themselves accessed by the dominant P-Demon as the P-Demon undergoes various states in which its sub-structures are either integrated into the dominant information web the P-Demon occupies or not<sup>171</sup>. Crucially, as these sub-structures are accessed, the dominant P-Demon would undergo a structural change (and thereby a phenomenal change) as it overlaps the various parts of the ‘information-web’ that these autocatalytic structures might form.

In this respect, the P-Demon might conceivably be undergoing a multiplicity of amplified experiential content simultaneously as it conceivably overlaps a multiplicity of these information webs simultaneously (in this sense, as shall be explored in the subsequent section, we might say that as the structure the P-Demon occupies becomes richer and more complex, so too does the phenomenal content the P-Demon undergoes). The potential

---

<sup>171</sup> In certain systems, we might even think of these lesser structures as P-Demons in their own right and think of the wider structure as the evolutionarily optimised outcome of a symbiotic relationship that forms between lesser P-Demons and more dominant P-Demons. Each P-Demon might employ its own capacity for phenomenal individuation in order to restructure itself so as to form a system that maximises the rapidity and accuracy of the informational entropy reduction relative to the experiences of the most dominant P-Demon, because doing this will maximise the chances of the dominant P-Demon accurately individuating between experiential types (and thereby maximise the chances of the system it occupies persisting in time), which facilitates more experiential content for each of these lesser P-Demon’s by virtue of increasing the entire system’s access to more experiential content (the various subsystems at work in the human body might therefore be thought of as manifestations of the work of these lesser P-Demons). In this sense, we might reframe Dennett’s (2016) sentiment that ‘we’re robots made of robots made of robots’ (p. 292) in panpsychist terms by suggesting that we’re P-Demons made of P-Demons made of P-Demons. The spatial properties at play remain continuous regardless of whether we think of our constituents as robots or P-Demons, the only difference is that the latter finds a place for phenomenality whereas the former does not.

consequences of this would be that 1) phenomenal experiences are amplified in complexity and richness and 2) phenomenal experiences become informationally laden. This is possible because when these structures are accessed they not only facilitate an experience that is itself richly structured (by virtue of containing a web of other phenomenal types integrated into a complex whole), but also potentially a structure that is itself rich in informational content by way of containing a phenomenal experience that is interconnected to other previously encountered phenomenal experiences in such a way that the informational entropy surrounding the given experience is conceivably, over the course of evolutionary optimisation of these structures, reduced.

To elucidate this, we might reconsider the image of the simplistic P-Demon attempting to differentiate cold particles from hot particles. If this P-Demon occupied a system with autocatalytic webs of information of the type I describe, it seems conceivable that upon undergoing enough 'hot' (1) and 'cold' (0) experiences, the P-Demon would evolve an information web pertaining to the experience of 1 that would facilitate the redefinition of 1 in terms of previous encounters with this particular phenomenal experience/stimulus, which would in turn allow for the demon to more rapidly delineate the differences between the experience of 1/0. In this example, the information web might contain previous experiences with '1', as well as previous structural states (and thereby the experiences they instantiate) that occurred upon '1' overlapping '0', or '0' overlapping '1' et cetera. In this sense, upon accessing a structure that facilitates the phenomenal experience '1', the P-Demon is also accessing an information web imbedded in this structure that amplifies the experience of '1', so that the P-Demon is accessing all of the previous structural states (and their accompanying 'what it is like' experiences) related to the structural state associated with '1' and thereby accessing all of the interactions with this particular experience and all of the previously

encountered phenomenal disparities/similarities between this experience and others of a different kind. In light of this, the amplified experience of '1' becomes indisputably informationally rich and complex as a result of being redefined in terms of previously encountered experiences in such a way that the dominant P-Demon now not only holds simplistic first-order knowledge pertaining to 'what it is like' to undergo a phenomenal change as state '1' overlaps state '0', it knows what it means to be in state '1' as distinct from state '0' by virtue of this amplified phenomenal experience being a part of a complexly interconnected map of the related phenomenal experiences that came before it. Over time, we might even conceive of how this process might become so refined and complex that the informational entropy surrounding any number of phenomenal experiences is reduced with such rapidity that the P-Demon might individuate the disparities between various experiential types (be those abstract or concrete) with increasing ease and efficacy, as Gabora (2002) elucidates:

*'As the density of abstractions increases, the number of recall paths amongst them increases exponentially; thus the probability that they crystallize into an interconnected network increases. Eventually the worldview becomes autocatalytically closed, in that it yields a potential pathway of associative recall from any one stored experience to any other. That is, for any stimulus that is encountered there exists a way of redefining that stimulus in terms of previously encountered stimuli.'* (p. 11)

If this remains uncontested, one might argue that as this process of amplifying phenomenal experiences becomes optimised, it is conceivable that the dominant P-Demon would evolve so as to sit atop a system constituted of various lesser (informational) structures, with each structure forming a new pathway of associative recall in order to facilitate the process of

rapid and complex phenomenal individuation. In this sense, the dominant P-Demon might conceivably, over time, utilise its capacity for increasingly rich and efficacious phenomenal individuation until a structure very much like the brain evolves<sup>172</sup>. With this in mind, I now turn to offer an insight into the structural mismatch problem that might shed some light on the relationship between the structure of the brain and the structure of our phenomenal experiences.

### 6.1.3 An insight into the Structural Mismatch Problem

The conscious experience you are currently undergoing is phenomenally ‘rich’ in so far as it is a structured experiential whole that is divided into various distinct phenomenal types. In this case, you are undergoing the phenomenal experience of symbols amalgamated on a page, and this experience is structured such that there is a distinct phenomenal type associated with the whiteness of the page, the blackness of the symbols, and even also the symbols themselves. Yet, the structure of this particular visual experience is only a facet of the wider experiential whole that one might be undergoing in this particular moment in space-time, for this visual experience is simultaneously accompanied by a multiplicity of other phenomenal types – the heat of one’s environment, the feel of one’s chair, the experiences of low informational entropy that accompany the experiences of these symbols as meaningful words and not mere squiggles on a page et cetera. This entire complexly structured phenomenon occurs as a seamless symphony of experiential content that is orchestrated and unified within a singular subject, and this subject is aware of its place amidst the symphony. The problem is that the structure of this complex experiential whole seems markedly distinct from the

---

<sup>172</sup> I hold that this is the case because conceivably the structures these P-Demons form will occupy myriads of variances until a variant is selected for survival, and therefore the fact that myriads of systems have survived in this actual world by virtue of evolving a particular brain-like variant is justification for this claim.

structure of the human brain, and as a consequence, it seems difficult to reconcile the structure of our consciousness with either the macro or micro level structure of the system that is said to house consciousness if, as my account of Russellian Micropsychism entails, the brain is purported to be the extrinsic nature of phenomenality to the extent that the structure of the brain ought to be isomorphic with the structure of one's phenomenal consciousness. In what follows, I suggest that the model I have outlined thus far might provide some insight into how we may begin to address this particular problem by employing the use of the concept of dominant phenomenal overlap and thinking of the brain as an evolutionarily optimised structure that facilitates phenomenal amplification and informational entropy reduction.

Chalmers (2016) elucidates the nature of this structural mismatch problem by highlighting the following seemingly inconsistent tetrad:

- 1) Microphenomenal structure is isomorphic to microphysical structure*
- (2) Microphenomenal structure constitutes macrophenomenal structure.*
- (3) Microphysical structure constitutes macrophysical structure.*
- (4) Macrophenomenal structure is not isomorphic to macrophysical structure. (p. 206)*

From here, Chalmers (2016) brings the nature of this tetrad and its consequences into focus by erected a more explicit argument aimed at rendering Constitutive Russellian Panpsychism false:

- 1) If Russellian panpsychism is true, microphenomenal structure is isomorphic to microphysical structure.*



*(2) If constitutive panpsychism is true, microphenomenal (and microphysical) structure constitutes macrophenomenal structure.*

*(3) Microphysical structure constitutes only macrophysical structure.*

*(4) If microphenomenal structure is isomorphic to microphysical structure, then any structure constituted by microphenomenal structure (and microphysical structure) is isomorphic to a structure constituted by microphysical structure.*

*(5) Macrophenomenal structure is not isomorphic to macrophysical structure.*

*(6) Constitutive Russellian panpsychism is false. (p. 208)*

In both cases, the point of dispute seems to clearly be the contentious nature of the claim that ‘macrophenomenal structure is not isomorphic to macrophysical structure’, and this appears especially pressing in the absence of a robust elucidation of precisely what is meant by the concepts of macrophenomenal structure and macrophysical structure. As Chalmers (2016) himself notes, it seems relatively reasonable to articulate a distinction between a ‘narrow’ account of macrophysical structure which points exclusively to the spatial-temporal properties described by physics, and a ‘broad’ account which points to ‘any structure constituted by microphysics: for example, chemical, biological, [informational], computational structure’ (p. 208). Again, as Chalmers points out, if this is the case then premise (3) appears suspect on the grounds that microphysical structure does not necessarily *only* constitute narrowly macrophysical structures, and therefore ‘one can suggest macrophenomenal structure is isomorphic to certain information structure in the brain’ (Chalmers 2016, p. 209).

Whilst Chalmers concludes that it is not immediately clear whether this particular line would work for the Russellian Micropsychist, I hold that by adopting a broad account of what

constitutes macrophysical structure we might employ the dominant phenomenal bonding solution to the combination problem advanced in the previous chapter, and the account of the potential evolution of P-Demons advanced in the previous sections, to arrive at one way in which we *can* conceive of macrophenomenal structure as being isomorphic to macrophysical structure.

On this account, as one instantiation of phenomenality comes to dominate others, its structure evolves in time to become richer and more complex until the most appropriate description of its nature is in terms of an informational structure in place to amplify phenomenality and thereby reduce informational entropy. In this sense, the macrophysical structure of the brain might be (broadly) described extrinsically as an informational structure<sup>173</sup> that has evolved to represent a multitude of experiential types and reduce the informational entropy associated with these experiential types to reveal the informational disparities between structured phenomenal content (by way of mapping incoming content to previous interactions with such content)<sup>174</sup>. The brain's macrophenomenal structure might therefore very well be isomorphic to the brain's macrophysical structure because the macrophysical structure of the brain has evolved to optimise the macrophenomenal structure of the dominant P-Demon. We might make sense of this by thinking of amplified instantiations of phenomenality as highly complex autocatalytic structures which, when overlapped by the dominant P-Demon, give rise to phenomenal content that has been appropriately amplified both in terms of the richness of its experiential content and the extent of the information that this content is laden with. In

---

<sup>173</sup> See Galadi et al (2019) for a recent attempt to frame the brain as an informational structure.

<sup>174</sup> In this respect the brain might conceivably represent a multiplicity of experiential types simultaneously by virtue of undergoing a multiplicity of structural states simultaneously, and each of these states might be unified within a singular experiential locus in which the informational entropy relative to these types is reduced to such an extent that this various 'types' become parts of 'what it is like' to be the dominant phenomenality that might be individuated as 'parts' by an analysis of the phenomenal and informational content associated with them.

this respect, brain states are informational structures that have evolved to represent amplified, rich phenomenal content of low informational entropy and optimise the P-Demon's capacity for phenomenal individuation in a way that will ensure the survival of the wider system<sup>175</sup>. And, if this is the case, the force of the structural mismatch problem appears to dissolve, for whilst the intrinsic nature of the brain might appear to be nonisomorphic to a narrow account of its extrinsic structure, this extrinsic structure might in fact be broadly construed as an informational structure replete with various interconnected sub-structures in place to represent precisely the sort of phenomenal content we undergo as the phenomenally conscious subjects that *are* the intrinsic natures of evolutionarily optimised brains<sup>176</sup>.

This model is not without issue, however, for it seems we are left with at least three discernible problems: 1) an account of why (and how) a dominant phenomenal structure occurred in the first place, 2) an account of why (and how) microphenomenal structure should give rise to broad macro phenomenal structure at all, and 3) an account of how broad macrophenomenal structures relate to narrow macrophenomenal structures, and an account of how broad macrophenomenal structures can stand in 'the constitutive relation that is plausibly required to avoid causal exclusion worries' (Chalmers 2016, p. 2010)<sup>177</sup>. With the problems articulated, I propose addressing (1) and (2) simultaneously by reasserting the

---

<sup>175</sup> We might also construe emotions as evolutionarily optimised phenomenal types derived from certain informational structures which serve to guide the P-Demon towards behaviours that are likely to perpetuate the P-Demon's persistence in time. We might therefore also suggest, as Lacalli (2020) does, that the evolution of this broadly informational structure 'shapes' experience by reducing the experiential or structural 'noise' that would preclude the P-Demon's capacity to accurately and efficiently individuate between the experiential content that is directly relevant to its survival.

<sup>176</sup> Further, because 'microphenomenal structure is isomorphic to microphysical structure, and any structure constituted by microphenomenal structure (and microphysical structure) is isomorphic to a structure constituted by microphysical structure' (Chalmers 2016, p. 206), as various brain structures are overlapped by the dominant phenomenality, the dominant phenomenality undergoes structured phenomenal content that is isomorphic to *both* a broad informational account of these macrophysical structure *and* a narrow account of macrophysical structure (because both are constituted by microphysical structure).

<sup>177</sup> I note that Chalmers (2016, p. 205-10) highlights the latter two of these problems.

sentiments espoused in the previous chapter, in which I highlight that whilst the dominant phenomenal bonding solution to the combination problem does not entail a robust account of *precisely why* this sort of bonding might occur, it does offer a potential solution to the combination problem that is not beyond the purview of liberal naturalism (and arguably therefore should be assessed on the extent of its relative explanatory power in relation to the combination problem when compared to the inadequacies of its counterparts, and not on the extent of its relative lack of explanatory power in this one area). With this in mind, I suggest that an account of why a broad informational structure might arise from microphenomenal structure is clear from the perspective of the evolution of dominant P-Demons, for if the Constructal account I have offered thus far is upheld, then in order for a system to persist in time it invariably restructures itself to provide greater access to the currents that run through it. It seems not unreasonable, then, to infer that in the case of human life the greatest access was achieved by the evolution of broad informational structures of the type we might associate with a broad account of the brain, and this would seem to satisfy an account of why (broad) macrophysical structures formed<sup>178</sup>. With this in mind, I now turn to address the final problem.

This problem seems to have occurred as a consequence of the inference that broad macrophysical structure exists in parallel to narrow macrophysical structure. In actuality, however, broad and narrow macrophysical structures are isomorphic because both are

---

<sup>178</sup> I concede that this in itself is not a sufficient articulation of precisely *how* this occurred and note that the specific details of such an account might coincide with advancements in evolutionary neuroscience (see Nani et al 2019 for a review of contemporary neuroscientific research into identifying the neural correlates of consciousness and a review of the problems associated with this line of research).

constituted by microphysical structure, and therefore if the argument for composition as identity I have previously offered holds, it is the microphysical structure taken as a whole that captures the identity of *both* broad and narrow macrophysical structure. Subsequently, whilst a narrow concept might capture macrophysical structure from the perspective of mass, space, time et cetera, and a broad concept might capture the macrophysical from the perspective of informational structures and amplified phenomenality, both concepts capture the same referent: the composition of microphysical structure into a brain-like whole. In this sense, there are no substantive causal exclusion concerns because the identity of the whole is one and the same as the parts of the whole *taken together*.

With this established, it does not appear unreasonable to contend that whilst the Russellian Micropsychist model I have articulated does not necessarily provide steadfast, indubitable solutions to the structural mismatch problem, it does provide some insight into the nature of the problem by way of offering an outline of the brain as an informational structure that has evolved to facilitate the fundamental disposition to phenomenally individuate (and subsequently optimise the faculties of P-Demons), and therefore I take it that I have offered a fitting account of why macrophysical structure might very well be isomorphic with macrophenomenal structure. With this as a foundation, I now draw my Constructal account of the evolution of phenomenality to a close by offering a theory of cognitive phenomenology which might explain the relationship between thought, awareness, phenomenality and P-Demons, and might subsequently serve to bring us one step closer to a complete micropsychist theory of consciousness.

## 6.2 The phenomenality of thought

If the sentiments I have espoused thus far are taken to be true, at this moment we might think of ourselves as P-Demons that are employing phenomenality to sustain a state that is far from thermodynamic equilibrium. If this is so, your evolutionarily optimised capacity to phenomenally individuate between the intricacies of this experiential content is what has guided you through this piece and is what is currently enabling you to delineate the meaning contained within these symbols. In this respect, the phenomenal content you are currently undergoing is content of low informational entropy relative to you (the dominant P-Demon), and, with this in mind, it might not appear overtly absurd to attest, in line with Strawson (1994) Pitt (2004), Horgan & Tienson (2002), Siewert (1998), Chudnoff (2015) and various others, that, in fact, the thoughts you are having pertaining to this piece are themselves also phenomenally constituted<sup>179</sup>, for they are not obviously different in kind to any other content – that is, when contextualised in terms of the theory I have presented, they are themselves also conceivably phenomenally constituted content that is of low informational entropy relative to the dominant P-Demon. In this context, then, thinking becomes the continual modulation of one's structure (and thereby one's experiential content) to expedite a system state in which the informational entropy relative to this content is lowered, and this might conceivably be achieved by a rapid process of evolutionarily optimised phenomenal individuation as the P-Demon employs thermodynamic work to restructure itself in order to manipulate the phenomenal content it is undergoing and facilitate the process of phenomenal individuation we reference as 'thinking'.

---

<sup>179</sup> 'As Strawson (1994) notes: the experience of seeing red and the experience of now seeming to understand this very sentence, and of thinking that nobody could have different parents...all fall into the vast category of experiential episodes that have a certain qualitative character for those who have them as they have them' (p. 194).

I therefore suggest that it is not immediately clear why thought, or cognition, ought to be exclusively understood as a non-phenomenal process<sup>180</sup>, and, with this said, I now turn to expand on this inference by offering a new take on the nature of cognitive phenomenology which might be employed to both sharpen our understanding of why there is ‘something it is like’ to consciously think/occupy cognitive states and perhaps even lead us to a solution to the awareness problem outlined within the introductory chapter of this piece.

### 6.2.1 A new take on Cognitive Phenomenology

*‘As we think – whether we are speaking in complete sentences, or fragments, speaking barely or not at all, silently or aloud – the phenomenal character of our noniconic thought is in continual modulation, which cannot be identified simply with changes in the phenomenal character of either vision or visualization, hearing or auroralization, etc*

(Siewert 1998, p. 282)

Siewert’s sentiments appear especially pertinent given the model I have elucidated thus far. To reiterate, on this account, ‘thought’ might be conceived of as simply phenomenal content of low informational entropy, such that the continual modulation of ‘thought’ becomes the continual modulation of one’s structure (and thereby one’s phenomenality), and this modulation is facilitated by an evolutionarily optimised process of informational entropy reduction that enables rapid, complex and efficacious phenomenal individuation between the informationally-rich content one is undergoing as this process of continual modulation occurs. If this is the case, it seems not unreasonable to posit that the continual modulation of

---

<sup>180</sup> In the sense that whilst there might very well be other (perhaps informational) theories of what constitutes thought, it does not hold that these ought to be considered as the *only* available accounts given that a phenomenality of thought (i.e., an account that explains thought in terms of phenomenality) might conceivably be advanced.

our phenomenal character *is* the continual modulation of our thoughts, and arguably adopting a form of cognitive phenomenalism of this kind would make sense not just of Siewert's sentiments as described above but also of the contemporary move to entwine thought with phenomenality<sup>181</sup>.

To comprehend how this might be the case (and how we might untangle the complexities of this so that it might be reconciled with our common-sense notions of what constitutes thought), in line with Chudnoff (2015) I begin by employing four examples to reveal how a change in our thoughts or cognition might be identified with a change in our phenomenal experience:

*[Understanding]. You are trying to read the instructions for a medicine a veterinarian prescribed for your dog. At first it is illegible. Then you see that it says to administer the medicine twice daily for one week.*

*[Intuiting] In a book you read, 'If  $a < 1$ , then  $2 - 2a > 0$ ,' and you wonder whether this is true. Then, you 'see' how  $a$ 's being less than 1 makes  $2a$  smaller than 2 and so  $2 - 2a$  greater than 0.*

*[Seeing] You are looking for your dogs in a dog park. At first you cannot pick them out of the mass of other dogs. But then you see them chasing a tennis ball.*

*[Reacting] In the news, you read about a factory building collapsing on Bangladeshi garment workers who were ordered to work despite warnings about the safety of the building. This makes you sad and angry' (Chudnoff 2015, p. 82).*

---

<sup>181</sup> See proponents of cognitive phenomenalism, e.g., Pitt (2004), Horgan & Tienson (2002), Siewert (1998), Chudnoff (2015).



In all four of the above cases, there is ‘something it is like [to undergo the system state that occurs] at the moment before understanding, intuiting, seeing or reacting, and there is something it is like [to undergo the system state that occurs] at the moment of understanding, intuiting, seeing or reacting – and what it is like before being in these mental states is different from what it is like while being in these mental states’ (Chudnoff 2015, p. 83). In this sense, it is difficult to argue against the claim that there is some phenomenality associated with cognition, for in all four cases there is a disparity between the phenomenal character that occurs before we comprehend some content and the phenomenal character that occurs upon our comprehension<sup>182</sup>.

While this may be the case, Chudnoff’s sentiments are, in this instance, employed to advance a general case for the reality of cognitive phenomenology, and therefore do not offer an explicit account of what might be happening in the transitional moments between our not comprehending some content and our comprehending it. In what follows, I intend to provide a new take on cognitive phenomenology that might shed some light on precisely what is going on during these moments. Subsequently, unlike Chudnoff (2015), whilst I do not devote any time to arguments in favour of the reality of cognitive phenomenology, I do take cognitive phenomenology as a postulate and attempt to broaden our understanding of this concept by likening the moments in which we come to understand, intuit, see, react, comprehend et cetera to moments in which we come to reduce informational entropy and phenomenally individuate with more accuracy. With this established, I see it fit to curtail the scope of this section to concern myself solely with an account of how the process of

---

<sup>182</sup> I note that Chudnoff’s concern is, in part, delineating whether there is a phenomenal disparity between sensory states in which we see or react and states in which we intuit or know. For my purposes, I do not devote any concern to this question for, as my model elucidates, it is not difficult to conceive of how all four instances might be construed as instantiations of phenomenally individuating between content that is of low informational entropy.

phenomenal individuation, and the account of the evolutionary optimisation of P-Demons I have offered thus far, might be exercised so to offer a new take on what is going on during the process of understanding, intuiting, reacting (for brevity henceforth referenced under the umbrella term ‘comprehending’). In this section I will therefore exclusively concern myself with the fundamental question: can a model be erected in which cognition (which I take to include thought and comprehension) might be construed as an evolutionarily optimised outcome of the disposition to phenomenally individuate?

The answer, as I have attempted to elucidate in this piece up until this point, is of course *yes*. In fact, given the account of phenomenal individuation on offer, I suggest that cognition is a wholly unsurprising outcome of the evolution of this disposition – indeed, one might even conceive of it as a necessary evolutionary outcome if, as I have intimated, what we reference as cognition is in fact structured phenomenality optimised by complex and efficient processes of informational entropy reduction and phenomenal amplification to minimize phenomenal surprise and expedite increasingly accurate and efficacious phenomenal individuation (and thereby facilitate more access to experiences and subsequently an increased chance of survival).

In this model, then, cognition is a consequence of the higher-order phenomenal structure that manifests in line with the evolutionary optimisation of P-Demons and is therefore at core an informational structure that facilitates the process of phenomenal individuation. If this is the case, we might conceive of the content of this structure as intrinsically phenomenal and suggest that this capacity for individuating between the content we are undergoing simply *is* the intrinsic nature of the processes we reference as cognition. In this respect, the various autocatalytic structures that are overlapped by the dominant instantiation of phenomenality

become the very informational structures that enable this dominant P-Demon to undergo phenomenal content which is of low informational entropy and thereby avoid high degrees of ‘phenomenal surprise’, and the process of individuation the P-Demon employs to select between the variances in this content just *is* the process that underpins what we reference as cognition.

To conceive of this, we might consider the activation of the various informational structures that occur within the structural state you are currently occupying. In this case, your structure is in a process of continual modulation as it encounters its environment. For each of these modulations, there is a corresponding informational structure that enables you, the purported dominant P-Demon, to undergo these various structural changes in such a way that they are appropriately phenomenally amplified in a manner that the content of your phenomenal conscious experience becomes a multimodal unified whole replete with various distinct phenomenal types (each of which are of low informational entropy relative to you). This is possible because the structure in which the dominant P-Demon manifests can conceivably undergo a variety of modulations simultaneously, and it has conceivably evolved to reduce the amount of informational entropy relative to each of these modulations in such a way that the modes in which these modulations are experientially presented vary based on the totality of their informational content. In this sense, from the inside, the various modulations that occur within a structure that has evolved to reduce the informational entropy pertaining to them might very feel very much like our cognitive phenomenal states feel as we move about our environment, delineate the disparities between various types of phenomenal content, and react appropriately. In other words, the processes that occur upon our coming to intuit, see, react, understand, or comprehend might very well be framed in terms of this unifying process of informational entropy reduction and phenomenal individuation, for the moment of

comprehension that occurs in each such case can be construed as a moment in which the informational entropy relative to some given phenomenal content is reduced to the extent necessary to enable more accurate processes of phenomenal individuation.

Reimagine, for example, the process by which we attempt to pick out our dog amidst a crowded park. We might describe this as a process in which, at the outset, the informational entropy relative to the phenomenal content we are undergoing is high. There are a great many dogs in the park, and therefore the possible informational states of this content are high. From this starting point, we begin reducing the informational entropy relative to our content by engaging in a process of phenomenal individuation in which we focus on specific phenomenal types and delineate whether it is, or is not, the type we seek. If it is, we have found our dog, and if it is not, we have still reduced the informational entropy residing in the phenomenal content taken as a whole, because the possible informational states of said content have been reduced. As a result, as we go about the process of reducing informational entropy in this way, we eventually reach a state in which we find that which we seek, and we achieve this by engaging in the process of phenomenal individuation and informational entropy reduction by individuating the phenomenal type we are seeing as one and the same as the memory of the phenomenal type associated with our dog. In this example, this entire cognitive process is arguably reducible to the very processes that might occur within an evolutionarily optimised P-Demon, for the P-Demon would have access to the structure that corresponds to the specific dog in question (which would, in turn, entail having access to phenomenal content of low informational entropy capturing the specific type of phenomenality we seek) and from here can employ its disposition for phenomenal individuation to delineate the similarities, or dissimilarities, between this type of phenomenality 'S' and the other types it undergoes as it occupies various '~S' structures that

represent its environment in various ways. At the moment in which the correct type is matched, the P-Demon ‘sees’ its dog, comprehends that this is a phenomenal type of the kind being sought, and reacts appropriately. This is, therefore, an entirely cognitive process captured ‘from the inside’ in terms of phenomenal individuation and phenomenal content of low informational entropy, and I propose that as this particular example acts as just one case of this process, it might not appear overtly radical to suggest that this same basic process might be sufficient to capture what is, at foundation, going on during all cognitive phenomenal states, for conceivably all such states might involve the basic tools of phenomenal individuation and structures of low informational entropy (and these tools might manifest in various intensities and complexities in various systems)<sup>183</sup>.

Whilst the above example might capture what is going on when we are interacting with ostensibly ‘concrete’ phenomenal types – that is, phenomenal types with a discernible external physical referent – it is perhaps a little less clear how the same might apply to abstract phenomenal types with no obvious physical referent. Further still, it is perhaps a little unclear how abstract phenomenal types are grounded at all. In this case, I propose, in line with Gabora’s work (2002), and my previous account of the optimisation of phenomenal individuation, that our phenomenal content relative to abstractions is not obviously different in kind to concretions, for both manifest as structural modulations that are, conceivably, optimised in terms of phenomenal amplification. Therefore, just as when we undergo phenomenal types pertaining to concrete referents various informational structures are accessed so that said phenomenal types are appropriately amplified in terms of structure and

---

<sup>183</sup> I note that a more sustained explanatory account would be required to convincingly substantiate this claim, and, as it stands, such an account is beyond the scope of what remains of this thesis. However, I also note that providing an account of cognitive phenomenology based on the model I have outlined might serve as an interesting focus for future research in this area.

informational content, so too are various informational structures accessed when undergoing phenomenal types pertaining to abstract referents. In fact, as the dominant P-Demon simply *is* an evolutionarily optimised informational structure, there is no obvious disparity (in so far as both are reduced to types of structural modulations) between undergoing phenomenal content pertaining to abstract referents or concrete referents. In both cases, the content is reducible to a structural modulation in the brain (or whichever informational structure the P-Demon occupies).

To extrapolate, I argue that much in the same way that informational structures form to facilitate our interactions with concrete referents, informational structures also form to capture the interactions *between* concrete referents, and it is these very informational structures that we reference as abstractions. In this case, as the wider informational structure we inhabit becomes increasingly optimised and complex, we develop new phenomenal types in line with the onset of new informational sub-structures (perhaps lesser autocatalytic structures) which exist as a conglomeration of the various disparities and/or similarities between the structures that represent concrete referents. And these complex, newly amplified, phenomenal types capture abstract referents of various kinds, which in turn facilitate new informational structures pertaining to the differences/similarities between these new abstract referents, and so on. In this sense, as the process of phenomenal individuation is supported by these various informational structures, the process of phenomenal individuation becomes increasingly optimised over time, which in turn facilitates more optimisation and more informational entropy reduction relative to more complex and convolved abstract phenomenal types. One might even argue that the onset of language stands as an example of this process, for language itself might be classified as series of complex phenomenal types that facilitate increased informational entropy reduction.

In this sense, we might even think of ourselves as having become evolutionarily optimised individuation machines working on the software of phenomenality and hold that every process we undergo is now a process which involves informational entropy reduction and phenomenal individuation in order to expedite more efficient manifestations of these processes. If this is the case, it appears that, over time, one consequence of these increasingly efficient processes might be that they are turned inward so that the P-Demon achieves awareness of its nature, and, with this said, I now turn to finalize this piece by addressing the awareness problem that was outlined at the outset.

### 6.2.2 Addressing the Awareness Problem

Before I elucidate how P-Demons might achieve self-awareness, it is prudent to qualify the discussion at hand by pointing out that, in this context, I am exclusively concerned with *reflective* self-awareness - that is knowledge that the P-Demon exists as a unified subject of experience replete with various dispositions, or perhaps the acknowledgement that this phenomenal subject 'S' is an object distinct from other objects '~S'. This qualification is necessary because, as intimated in the previous chapters, phenomenal properties arguably entail a minimal subject that is, in some sense, immediately and pre-reflectively 'aware' of its own consciousness by virtue of holding first-order knowledge pertaining to the feel of what it is like to be this instantiation of consciousness. Consequently, this section is concerned with delineating how we move from this pre-reflective self-awareness to a reflective state of awareness in which the subject is aware not just of its immediate phenomenality, but also of itself as a phenomenal subject replete with various dispositions and qualities.

I therefore hold that the task at hand should not be overly difficult given the account of the dispositions of P-Demons (and the potential evolutionary outcomes of said dispositions) outlined in the previous sections. This is because, arguably, if the sort of awareness we require is simply the P-Demon attaining the capacity to individuate itself as something distinct from its environment, this should pose no challenge at all given that P-Demons are disposed to this sort of individuation, and the evolutionary optimisation of this disposition is entirely likely to facilitate a process of informational entropy reduction that leads to the P-Demon demarcating itself as something distinct from its environment.

We might even say that as this process of informational entropy reduction continues in time, the phenomenal map of oneself and one's nature invariably becomes richer in line with one's phenomenal map of one's environment. This is because the P-Demon, when optimised, is an individuating machine that delineates the disparities/similarities between its content with increasing rapidity and ease, and so if a set of content seems to share a similarity, the P-Demon would, if optimised in terms of selection, eventually exhibit the sort of informational structure necessary in order to spot this. In this case, if all of the content the P-Demon undergoes shares the quality of inciting a phenomenal feel, or having a phenomenal character, then the P-Demon would evolve an informational structure that captures the similarities between this specific unifying quality, and, as the process of informational entropy reduction and phenomenal individuation becomes increasingly optimised, perhaps eventually surmise that as each of these seemingly disparate instantiations of phenomenal content share this unifying quality, there must be something that unifies all of this content: the subject itself. In this sense, the map of one's environment facilitates the map of oneself (and vice versa), for the P-Demon learns of objects that do not seem to share this unifying quality, and eventually forms informational structures pertaining to those objects (~S) that are



distinct in kind from the P-Demon (S). At this time, the P-Demon might very well be described as reflectively self-aware, for it separates itself from its environment by way of individuating the disparities between itself as a seer and that which it sees.

Further still, this same process might be employed to achieve awareness of its various dispositions, for each instantiation of experiential content is presented to the P-Demon in various ways that reveal several necessary qualities and dispositions that must manifest for this experiential content to present in the way it does. In this respect, it is not overtly difficult to conceive of how the P-Demon might employ its evolutionarily optimised capacity for phenomenal individuation and informational entropy reduction to delineate the disparities between dispositions and qualities that appear to be necessary for this content to present as it does and those that do *not* appear to be essential to this content. And we might conceive of how this map of these various necessary qualities/dispositions might eventually become so complex that the dominant phenomenality eventually achieves awareness that it is itself the bed of the various qualities and dispositions which are necessary for experiential content to manifest in the manner that it does. Subsequently, therefore, we might even conceive of a certain instantiation of these various dispositions and qualities employing its awareness of its existence, and its capacity for optimised phenomenal individuation, to advance a theory in which it posits itself as a P-Demon.

Perhaps unavoidably, however, the P-Demon might struggle to individuate its nature in totality and as a result we construct theses such as these in which we continue to claw at our true nature to unravel the fundamental question: what are we? As we reach the end of this thesis, I hope that I have at least offered a fitting *attempt* to address this question and hope that, perhaps, we are now able to offer a more fitting answer than we were at the outset. With

this said, I now turn to conclude this piece by recapping the nature of this undertaking and addressing some of the latent inadequacies in a theory of consciousness of the type I have endeavoured to construct.

## Chapter 7

### Conclusion

If we were to reimagine the infinity dice outlined within the opening of this piece, we might now, buoyed by the adoption of Russellian Micropsychism and the account of consciousness I have offered, be able to account for precisely how the rolling of these dice might contain the possibility of phenomenally conscious subjects replete with the cognitive faculties that appear necessary in order to conceive of abstract concepts like infinity dice. In this respect, this piece has pursued a lofty goal – a ‘complete micropsychist theory of consciousness’ - that begins with the chaotic, minimal phenomenality that manifests at the bottom-level and has endeavoured to offer a framework which explains how this bottom-level phenomenality might transition into fully-fledged conscious subjects of the type we identify with.

This work has pursued this lofty goal by attempting to paint new colour into our understanding of consciousness from the perspective of Russellian Micropsychism, and it goes about this task by employing the core ingredients of phenomenality, phenomenal individuation, and informational entropy reduction – each of which are entailed by the broader concept of evolutionarily optimised ‘P-Demons’, and each of which has its own place in the picture of consciousness I have attempted to construct. In this respect, the picture I have drawn has been built upon a series of more rudimentary and uncontroversial propositions in the hope that even if the conclusions I have ultimately reached are not true, they are well-conceived. By means of recapping precisely how the picture of consciousness we are left with has been constructed (and familiarizing ourselves with these purportedly uncontroversial propositions), I now turn to offer a brief outline of the presuppositions that have led us from minimal phenomenality to P-Demons.

At the outset of the piece, I espoused the virtues of Russellian Micropsychism and attempted to offer a new take on the nature of phenomenality which turned on the notion that phenomenal properties are powerful qualities that instantiate structured experiences-for and experiences-of disposed to all the manifestations observed in physics *and*, when standing in the right reciprocal relations, disposed to phenomenally individuate between the structured states they occupy. From here, I built upon this model to provide an account of how we might find a place and role for phenomenality in the natural world, and suggested that the proposed bridge between the sort of minimal phenomenality manifesting at the lower-level and the sort of phenomenality manifesting in higher-order conscious subjects might be found in the concept of the 'P-Demon' - that is, the naturalised adaptation of a Maxwellian Demon which might employ phenomenal individuation to sustain thermodynamic disequilibrium, and eventually perhaps achieve the evolutionary optimisation necessary to turn this capacity for phenomenal individuation to more complex and lofty pursuits. With this in place, I offered a series of solutions to the combination problems which culminated in the possibility of a dominant instantiation of phenomenality that has parts and is perhaps able to, when standing in the right reciprocal relations, employ phenomenal individuation to sustain its access to new experiential states (i.e., act like a 'P-Demon'). The final passages of this thesis, therefore, turned on the notion that we, at this moment, might be P-Demons employing the tools of evolutionarily optimised informational entropy reduction and phenomenal individuation to delineate the veracity of the sentiments found herein.

At the close, I recognise that this is undeniably a relatively radical claim, yet I suggest that our initial concerns ought to be curtailed for construing our nature as such leads us to a neat model of consciousness which provides a means of reconciling bottom-level powerful

phenomenal qualities with their manifestations at higher levels, is wholly (liberally) naturalizable and is not inconsistent with the nature of phenomenality (and cognition) as it is actually experienced in higher level systems. In this respect, I hold that whilst such a conception might very well be radical, it is not unreasonable. For if informational entropy reduction is not the process by which we delineate the veracity of a sentiment, we must ask how one formulates their denial (or doubt) without employing some capacity for informational entropy reduction and thereby self-stultifying. Further, if this informational entropy reduction does not occur within a unified and bounded phenomenal locus that employs at least some capacity for the phenomenal individuation that appears necessary to reasonably ascribe the doubt to a knowing subject capable of doubting, we must ask from where this capacity originates, and more: for what purpose, if the origination of said doubt is not phenomenal in nature, would one be motivated to doubt such things at all? In this sense, in closing we might finally tentatively take a stand on the puzzle posed by Hawkins (1989, p. 174): it is phenomenality that breathes fire into the equations, ‘for everything begins with consciousness and nothing is worth anything except through it’ (Camus 1955, p. 10).

With this said, the description I have offered of human consciousness and cognition is not without issue. Indeed, this entire thesis, as the title suggests, ought to be thought of as just a step *towards* a complete micropsychist theory of consciousness, for, not least because of the scope of the undertaking, I have arguably not adequately addressed concerns involving attention, the subconscious, sleep, or the unconscious, nor have I adequately addressed the structural mismatch problem, or offered a full, empirically verified, and substantive, solution to the standard combination problem. Whilst the latter two of these not inconsiderable issues are problems that I have addressed elsewhere and take to be beyond the scope of this thesis, I briefly address the former of these issues in what follows.

In the case of attention, I suggest that as the dominant P-Demon employs its capacity for phenomenal individuation to dictate the flow of thermodynamic work in a system, the attention of this dominant instantiation of phenomenality is relative to the extent of the work being performed. In this case, if the flow of energy within the system is directed towards a given informational structure, this informational structure will occupy most of the attention of the P-Demon. Similarly, in the case of subconscious processing, we might think of this as the sort of restructuring that occurs even in the absence of the consciousness of the P-Demon, such that those informational structures which are not currently being overlapped by the dominant instantiation of phenomenality are still undergoing processes of restructuring in order to further reduce informational entropy and optimise the amplification of phenomenality that occurs upon their being overlapped. These processes are what we might think of as the subconscious. Further, sleep and unconsciousness might be conceived of as system states in which the amount of total phenomenal overlap is either diminished or reduced entirely in line with reductions to the number of connections being formed in the structure, to the extent that as the connections within the structure the dominant instantiation of phenomenality occupies are diminished in size, frequency of scope, so too is the extent of one's consciousness. In this sense, as one sleeps, the neuronal activation and synaptic connections in the brain are diminished, and conceivably so too is the extent of phenomenal overlap and thereby the extent of phenomenal consciousness. What we reference as unconsciousness, then, is simply a further reduction in the extent and complexity of the connections that occur in the structure of consciousness, to the extent that as more simplistic structures are occupied, phenomenality is reduced back to increasingly base states until we reach its most fundamental and simplistic state (i.e., a state of minimal phenomenality of the type associated with bottom-level instantiations of physicality). With these problems (briefly)

addressed, I reiterate that whilst I have attempted to offer some insight into the link between phenomenality and cognition, I do not take this thesis to have offered a full elucidation of the precise nature of cognition, nor do I take it to have offered a substantive account of the inner workings of cognition. In this respect, whilst the theory I have offered does at least seem to shed light on some of the more standard concerns surrounding theories of consciousness, I note these explanations are not without fault and that a full and robust articulation and defence of each of these accounts would require a more focused and sustained line of argumentation and research than that which I can offer in what remains of this thesis.

With this said, I hold that we ought to judge the success of our metaphysical models of reality on the extent to which said models purport to provide solutions, and in this case the solutions offered, whilst speculative, are potentially vast. For example, a micropsychist theory of consciousness of the type I have endeavoured to advance solves the ‘hard’ problem, finds an explicit place and role for phenomenality and explains the perennial mystery: *what are we?* Further, adopting this view potentially leads to an enriched understanding in other areas: it might provide a new account of the foundations of logic, for it is not especially difficult to conceive of how individuating the disparities/similarities between content might potentially, over the course of its evolutionary optimisation, lead us to a process of individuation and informational entropy reduction that is not beyond reconciliation with the law of contradiction, the law of excluded middle, or the law of identity; it might provide new insights into the nature of knowledge, for knowledge would become merely a process of reducing the informational entropy of phenomenal content to states of 0 or near 0<sup>184</sup>; it might

---

<sup>184</sup> One might even suggest that the reason a priori knowledge is often considered more epistemically robust than a posteriori knowledge is because our capacity for informational entropy reduction is not limited for the former in the same way it is for the latter. To simplify, we might never achieve a state of absolute 0 in which all

explain the move towards increased technological and theoretical advancements, for the process of optimizing phenomenal individuation and informational entropy reduction might very well flow in the direction of increases in both the rapidity, complexity and accuracy of these processes *and* reductions to the local entropy produced during these processes (or, put differently, the amount of work involved in the facilitation of these processes). Such that as these processes become increasingly optimized, the P-Demon might conceivably move in the direction of states that serve to reduce the local amount of energy expended (and thereby reduce the extent of the local entropy produced relative to the P-Demon) in the process of surviving and experiencing. The motivation to continue innovating might therefore be one born of a motivation to optimize our capacity as a P-Demon by embedding ourselves in an environment in which the process of individuation faces less resistance (or is made more efficient or 'easier') by way of being directed towards the production of technologies that result in reductions to both the amount of informational and physical entropy that accompany the process of continuing to experience and persist in time. Finally, it might even explain the nature of philosophising itself, for our being drawn to such pursuits is itself explained by the P-Demon's evolved tendency to restructure itself such that it might reduce informational entropy not just more rapidly and efficaciously, but also more imaginatively and accurately. With all of this said, one might even suggest that a theory of P-Demons perhaps even explains itself, for, if P-Demons exist, a naturalised theory of metaphysics that incorporates them explains its own existence by way of the theory itself being a likely outcome of the reality of P-Demons (and the dispositions they might hold). In this sense, what is left? Either a denial of the reality of one's lived phenomenal experience (or its causal efficacy), or the

---

possible informational states pertaining to empirical data are revealed because we do not hold the necessary epistemic tools in order to indubitably demarcate all of the potential informational states pertaining to such data, but the scope of our enquiry (and thereby the possible informational states pertaining to the content at hand) is lessened and curtailed dramatically in the case of a priori knowledge, and therefore achieving such a state is perhaps less troublesome.



attestation that phenomenality is a causally efficacious, non-illusory phenomenon that has a tangible role in the natural world, and subsequently therefore perhaps some openness to the possibility that you yourself are a P-Demon.

*Word count: 99936 excluding references/tables.*

## References

- Aad, G et al (ATLAS collaboration). (2015) *Search for new phenomena in the dijet mass distribution pp using collision data at  $\sqrt{s} = 8$  TeV with the ATLAS detector*. Physical Review D, 91 (5).
- Acuna, P & Dieks, D. (2014) *Another look at empirical equivalence and underdetermination of theory choice*. European Journal for Philosophy of Science 4: 153-180
- Allen, S. (2012) "*What Matters in (Naturalized) Metaphysics?*". Essays in Philosophy 13 (1): 211-241
- Armstrong, D. (1997). *A World of States of Affairs*. Cambridge University Press.
- Bailey, A. (2006) "*Zombies, Epiphenomenalism, and Physicalist Theories of Consciousness*". Canadian Journal of Philosophy 36: 481–509.
- Barge, L et al (2017). *Thermodynamics, Disequilibrium, Evolution: Far-from-equilibrium Geological and Chemical considerations for Origin-of-Life Research*. Origin of Life and Evolution of Biospheres 47 (39-56).
- Basile, P. (2010) *It must be True – But How Can It Be? Some Remarks on Panpsychism and Mental Composition*, in Basile, P, Kiverstein, J and Phemister, P. (2010) *The Metaphysics of Consciousness Royal Institute of Philosophy Supplement: 67*. Cambridge University Press
- Baxter, D. (1988) Identity in the Loose and Popular Sense. Mind 97: 576-582
- Bayne, T. (2001) *Co-Consciousness Review of Barry Dainton's 'Stream of Consciousness'*. Journal of Consciousness Studies 8 (3)
- Bejan, A & Lorente, S. (2011) *The Constructal Law and the Evolution of Design in Nature*. Physics of Life Review 8 (3), 209-240

- Bejan, A. (1997) *Constructal Theory Network of Conducting Paths For Cooling a Heat Generation Volume*. International Journal of heat and mass transfer 40 (4): 799-816
- Bejan, A. (2005) *The Constructal Law of Organisation in Nature: Tree-Shaped Flows and Body Size*. Journal of Experimental Biology 208: 1677-1686
- Bejan, A. (2010) *The Constructal Law of Design and Evolution in Nature*. Philosophical Transactions of the Royal Society B 365 (1545): 1335–1347
- Bejan, A. (2012) *Design in Nature: How the Constructal Law Governs Evolution in Biology, Physics, Technology and Social Organisation*. DoubleDay.
- Bejan, A. (2014) *Maxwell's Demons Everywhere: Evolving Design as the Arrow of Time*. Scientific Reports (4): 1-4.
- Bejan, A. (2016) *Life and Evolution as Physics*. Communicative & Integrative Biology 9 (3): 1-13.
- Bejan, A. (2016) *The Physics of Life: The Evolution of Everything*. Macmillan USA.
- Bennett, C. (1982). "The thermodynamics of computation—a review". International Journal of Theoretical Physics 21(12): 905–940
- Bergami, M & Berninger, B. (2012) *A Fight for Survival: The Challenges Faced by a Newborn Neuron Integrating in the Adult Hippocampus*. Developmental Neurobiology 72 (7): 1016- 1031.
- Berkeley, G. (1710/1982) *A Treatise Concerning the Principles of Human Knowledge*. Hackett.
- Bitbol, M. (2007) *How is emergence possible?*, in Auletta, G. (2007) *The controversial relationship between science and philosophy: a critical assessment*. Libreria Editrice Vaticana.
- Blum, P. (2018) *Existence and Essence*. The Analogy of Being Workshop, Ligerz. Accessed from [http://philipp.philosophie.ch/handouts/2018-3-1-essence\\_existence.pdf](http://philipp.philosophie.ch/handouts/2018-3-1-essence_existence.pdf) on 08/10/2020.

- Bohm, D. (2005) *Wholness and the Implicate Order*. Routledge.
- Bradley, H. (1893/2003) *Appearance and Reality: A Metaphysical Essay*. Adamant Media Corporation
- Brewer, B. (2006) *Perception and Content*. *European Journal of Philosophy* 14 (2): 165-181.
- Brody, A. (1980) *Identity and Essence*. Princeton University Press.
- Brüntrup, G. (2016) *Emergent Panpsychism*, in Brüntrup, G & Ila, L. (2016) *Panpsychism Contemporary Perspectives*. Oxford University Press.
- Bunge, M. (1982) The Revival of Causality. *Contemporary Philosophy* 2: 133-155
- Busby, C. & Howard, C. (2017) *Re-Thinking Biology—I. Maxwell's Demon and the Spontaneous Origin of Life*. *Advances in Biological Chemistry* 7: 170-181
- Byrne, A. (2004) *What Phenomenal Consciousness is Like*, in Gennaro, R. (2004) *Higher-Order Theories of Consciousness: An Anthology*. John Benjamins Publishing Company.
- Camus, A. (1955) *The Myth of Sisyphus and Other Essays*. Vintage Books: New York.
- Carruthers, P. (2005) *Consciousness: Essays from a Higher-Order Perspective*. Oxford University Press.
- Caruana, L. (2000) *Holism and the Understanding of Science*. Routledge.
- Chalmers, D. (1995) *Facing up to the problem of consciousness*. *Journal of Consciousness Studies* 2 (3): 200-19
- Chalmers, D. (1996) *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, D. (2004) *The Representational Character of Experience*, In Leiter, B. (2004) *The Future of Philosophy*. Oxford University Press. Accessed online from <http://consc.net/papers/representation.pdf> 25/07/19
- Chalmers, D. (2010) *The Character of Consciousness*. Oxford University Press.

- Chalmers, D. (2013) “*Panpsychism and Panprotopsychism.*” The Amherst Lecture in Philosophy 8: 1–35.
- Chalmers, D. (2016) *Panpsychism and Panprotopsychism*, in Bruntrup, G & Jaskolla, L. (2016) *Panpsychism: Contemporary Perspectives*. Oxford University Press.
- Chalmers, D. (2016) *The Combination Problem for Panpsychism*, in Bruntrup, G & Jaskolla, L. (2016) *Panpsychism: Contemporary Perspectives*. Oxford University Press.
- Chalmers, D. (2018) *Idealism and the Mind-Body Problem*, forthcoming in Seager, W. (2018) *The Routledge Guide to Panpsychism*. Routledge. Accessed from <https://philpapers.org/archive/CHAIAT-11.pdf> on 26/03/20 21
- Chalmers, D. (2018) The Meta-Problem of Consciousness. *Journal of Consciousness Studies*, 25 (9-10): 6-61
- Chibeni, S. (2004) *Holism in Microphysics*. *Epistemologia* 27 (2): 227-44.
- Chudnoff, E. (2015) *Phenomenal Contrast arguments for Cognitive Phenomenology*. *Philosophy and Phenomenological Research*, 91 (1): 82 – 104.
- Chudnovsky, E. (1985) *Thermodynamics of Natural Selection*. *Journal of Statistical Physics*, 41: 877-897.
- Coleman, S. (2009) *Mind Under Matter*, in Skrbina, D. (2009) *Mind That Abides: Panpsychism in the New Millennium*. Philadelphia: John Benjamins: 83-108.
- Coleman, S. (2013) *The Real Combination Problem: Panpsychism, Micro-Subjects, and Emergence*. *Erkenntnis* (1): 1 -26
- Coleman, S. (2014) *The Real Combination Problem: Panpsychism, Micro-Subjects, and Emergence*. *Erkenntnis* 79 (1): 19-44.
- Coleman, S. (2014) The real combination problem: Panpsychism, micro-subjects, and emergence. *Erkenntnis*, 79 (1): 19–44

- Coleman, S. (2016) *Panpsychism and Neutral Monism: How to Make up One's Mind*, in  
 Brüntrup, G & Jaskolla, L. (2016) *Panpsychism Contemporary Perspectives*. Oxford  
 University Press.
- Correia, F & Skiles, A. (2017) *Grounding, Essence, and Identity*. Philosophy and  
 Phenomenological Research 98 (3), 642-670
- Cotnoir, A. (2014) *Composition as Identity – Framing the Debate*, in Cotnoir, A & Baxter, D.  
 (2014) *Composition as Identity*. Oxford University Press.
- Crane, T. (2009) *Is Perception A Propositional Attitude*. The Philosophical Quarterly 59 (236):  
 452-469.
- Dainton, B. (2000) *Stream of Consciousness: Unity and Continuity in Conscious  
 Experience*. Routledge.
- Darwin, C & Wallace, A. (1858) 'On the Tendency of Species to Form Varieties; and on the  
 Perpetuation of Varieties and Species by Natural Means of Selection. Journal of the  
 Proceedings of the Linnean Society of London, Zoology 3 (9): 45-62.
- De Brigard, F. (2014) *In defense of the self-stultification objection*. Journal of Consciousness  
 Studies, 21 (5-6): 120-13
- Dennett, D. (1991) *Consciousness Explained*. Back Bay Books.
- Dennett, D. (1993) *Consciousness Explained*. Oxford University Press
- Dennett, D. (2016) in Skillen, J. (2019) *God's Sabbath with Creation: Vocations fulfilled, and  
 Glory unveiled*. Wipf and Stock Publishers.
- Dennett, D. (2018) *Facing up to the Hard Question of Consciousness*. Philosophical Transactions  
 of the Royal Society B 373 (1755): 20170342.
- Descartes, R. (1641/1996) *Meditations on First Philosophy*. Cambridge University Press.
- Dessalles, J & Zalla, T. (2011) *On the Evolution of Phenomenal Consciousness*. Quantitative  
 Biology, accessed from <https://arxiv.org/abs/1108.4296> on 03/03/2020.

- Dietrich, E. (2011) *There is No Progress in Philosophy*. *Essays in Philosophy* 12 (2): 9.
- Dipert, R. (1997) The Mathematical Structure of the World: The World as Graph. *The Journal of Philosophy* (94) 7: 329-358
- Eddington, A. (1927/2019) *The nature of the physical world: The Grifford Lectures 1927*. Klaus-Dieter Sedlacek.
- Elitzur, A. (1994) *Let There Be Life: Thermodynamic Reflections on Biogenesis and Evolution*. *Journal of Theoretical Biology* 168 (4): 429-459
- Esfeld, M. (2004) *Quantum Entanglement and a Metaphysics of Relations*. *Studies in History and Philosophy of Science part B: Studies in History and Philosophy of Modern Physics*, 35 (4): 601-617
- Feldman, H & Friston, K. (2010). *Attention, uncertainty, and free-energy*. *Front. Hum. Neurosci.* 4:215
- Fine, K. (2001). *The question of realism*. *Philosophers' Imprint* 1: 1–30.
- Fine, K. (2012a). *Guide to ground*, in Correia, F, & Schnieder, B. (2012) *Metaphysical grounding: Understanding the structure of reality*. Cambridge University Press
- Fowler, T. (2007) *The Evolution Controversy: A Survey of Competing Theories*. Baker academic.
- Fox, T. (2009) *Why Quarks are Unobservable*. *Philosophia Scientiae* 13 (2): 167-189
- Francescotti, R. (1999) *How to Define Intrinsic Properties*. *Noûs*, 33: 590–60
- Frankish, K. (2016) *Illusionism as a Theory of Consciousness*. *Journal of Consciousness Studies*, 23 (11-12): 11-39.
- Frankish, K. (2016) *Why Panpsychism is Probably Wrong*. Accessed from <https://www.theatlantic.com/science/archive/2016/09/panpsychism-is-wrong/500774/> on 03/03/2020.

- Friston, K. (2009). *The free-energy principle: a rough guide to the brain*. Trends Cogn. Sci 13: 293–301
- Friston, K. (2012). *A free energy principle for biological systems*. Entropy 14: 2100–2121.
- Fuchs, C. (2003) *Co-Operation and Self-Organisation*. TripleC 1 (1), 1-52.
- Gabora, L. (2002) Amplifying phenomenal information: Toward a fundamental theory of consciousness. *Journal of Consciousness Studies* 9 (8): 3-29
- Galadi, J., Pereira, S., Perl, Y., Kingelbach, K., Gayte, I., Laufs, H., Tagliazucchi, E., Langa, J & Deco, G. (2019) *Dynamical Informational Structures characterize the different brain states of wakefulness and deep sleep*. BioRxiv, 84667.
- Gertler, B. (2019) *Acquaintance, Parsimony and Epiphenomenalism*, in Coleman, S. (2019) *The knowledge argument*. Cambridge University Press.
- Giannotti, J. (2019) *The Identity Theory of Powers Revised*. Erkenn. Accessed from <https://doi.org/10.1007/s10670-019-00122-5> (08/08/2020)
- Gibran, B. (2014) *Causal Realism in the Philosophy of Mind*. Essays in Philosophy 5 (15) 2: 299-313
- Gleick, J. (2012) *Information: A History, a theory, a flood*. Fourth Estate.
- Goff, P. (2009) ‘Can the Panpsychist get around the combination problem?’, in Skrbina, D. (2009) *Mind that abides: Panpsychism in the New Millennium*. Benjamins Publishing.
- Goff, .P. (2009) Why Panpsychism doesn’t help explain Consciousness. *Dialectica*, 63.
- Goff, P. (2015) ‘*Against Constitutive Russellian Monism.*’, in Alter, T & Nagasawa, Y. (2015) *Consciousness in the Physical World: Perspectives on Russellian Monism*. Oxford University Press
- Goff, P. (2015) *Against Constitutive Panpsychism*, in Alter, T & Nagasawa, Y. (2015) *Consciousness in the Physical World*. Oxford University Press.



- Goff, P. (2016) *The Phenomenal Bonding Solution to the Combination Problem*, in Jaskolla, L. (2016) *Panpsychism: Contemporary Perspectives*. Oxford University Press.
- Goff, P. (2017) *Consciousness and Fundamental Reality*. Oxford University Press
- Goff, P. (2017) *Cosmopsychism*, in Schneider, S & Velmans, M. (2017) *The Blackwell Companion to Consciousness*. Wiley Blackwell Publishing.
- Goff, P. (2018) *Cosmopsychism, Micropsychism, and the Grounding Relation*, forthcoming in Seager, W. (2018) *The Routledge Guide to Panpsychism*. Routledge.
- Goff, P. (2019) *Did the Universe Design Itself?* International Journal for Philosophy of Religion 85 (1): 99-122.
- Goff, P. (Forthcoming 2018) *Cosmopsychism, Micropsychism and the Grounding Relation*, in Seager, W. (Forthcoming 2018) *The Routledge guide to Panpsychism*. Routledge.  
 Accessed from <https://philarchive.org/archive/GOFCMA> on 15/08/2018.
- Gottfredson, L. (1997) *Mainstream Science on Intelligence: An editorial with 52 signatories, history and bibliography*. Intelligence 24: 13-23
- Guillot, M. (2017) *I Me Mine: On a confusion concerning the subjective character of experience*. Review of Philosophy and Psychology 8: 23-53
- Gustaffson, J. (2011) *Phenomenal Continuity and the Bridge Problem*. Philosophia 39 (2): 289-296
- Hanna, R & Maiese, M. (2009) *Embodied Minds in Action*. Oxford University Press.
- Harnad, S. (1990) The Symbol Grounding Problem. Physica D: Nonlinear Phenomena, 42 (3): 335-346
- Hawking, S. (1989) *A Brief History of Time: From the Big Bang to Black Holes*. Bantam.
- Heard, D. (2005) *A New Problem of Ontological Emergence*. The Philosophical Quarterly 56 (222): 55-62
- Heidegger, M. (1927/1978) *Being and Time*. Wiley-Blackwell.

- Heil, J & Robb, D. (2003) *Mental Properties*. American Philosophical Quarterly 40 (3): 175-196.
- Heil, J. (2003). *From an Ontological Point of View*. Oxford: Oxford University Press
- Heil, J. (2012) *The Universe As We Find It*. Oxford University Press.
- Heil, J. (2015) *The universe as we find it*. Oxford University Press.
- Hibbert, B. (2018) *Against the physicalist zeitgeist: evaluating the contemporary case for physicalism and the prospect of property dualism*. Masters Thesis, Keele University.
- Horgan, T & Tienson, J. (2002) *The Phenomenology of Intentionality and the Intentionality of Phenomenology*, in Chalmers, D. (2002) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press.
- Hornsby, J. (2005). *Truth without truthmaking entities*, In Beebe, H & Dodd, J. (2005) *Truthmakers: The Contemporary Debate*. Clarendon Press, pp. 33–47
- Hornsby, J. (2005). *Truth without Truthmaking Entities*, in Beebe, H & Dodd, J (2005) *Truthmakers: The Contemporary Debate*. Clarendon Press, pp. 33–47
- Howell, R. (2014) *The Russellian Monist's Problems with Mental Causation*. The Philosophical Quarterly 65 (258): 22–39
- Hume, D. (1740/2012) *A Treatise of Human Nature*. Penguin Classics
- Humphrey, N. (1992/1999) *A History of the Mind*. Springer.
- Husserl, E. (1931/2012) *Ideas: General Introduction to Pure Phenomenology*. Routledge Classics.
- Huxley, T. H., 1874, “*On the Hypothesis that Animals are Automata, and its History*”, The Fortnightly Review 16 (New Series): 555–580.
- Jacobs, J. (2011) *Powerful Qualities, Not Pure Powers*. The Monist 94 (1): 81-102
- James, W. (1890/1950) *The Principles of Psychology*. Henry Holt.

Jaworski, W. (2016). *Structure and the metaphysics of mind*. New York: Oxford University Press.

Jia, T; Chandru, K; Hongo, Y; Afrin, R; Usui, T; Myojo, K & Cleaves, J. (2019) *Membraneless Polyester Microdroplets as Primordial Components at the Origins of Life*. PNAS 116 (32).

Kadic, N. (2017) *The grounding problem for panpsychism and the identity theory of powers*. Croatian Journal of Philosophy 17 (1):45-55

Kadic, N. (2020) *Dynamic Selves: An Outline for a New Panpsychism*. The University of Arizona: The Science of Consciousness Conference 09/18/20.

Kaila, V & Annala, A. (2008) *Evolution as Described by the Second Law of Thermodynamics*, accessed from <https://phys.org/news/2008-08-evolution-law-thermodynamics.html> (22/07/2021).

Kaila, V. & Annala, A. (2008) *Natural Selection for Least Action*. Proceedings of the Royal Society A, 464 (2099).

Kastrup, B. (2018) *The Universe in Consciousness*. Journal of Consciousness Studies, 25, No. 5–6: 125–55

Kauffman, S. (1993) *Origins of Order*. Oxford University Press

Kim, J. (1998). *Mind in a Physical World*. Boston MA: MIT Press.

Kim, J. (2005). *Physicalism, or Something Near Enough*. Princeton University Press

Koch, C. (2017), in Reardon, S. (2017) *A Giant Neuron Found Wrapped around Entire Mouse Brain*. Nature 543: 14-15

Koch, C & Crick, C. (2005) *What is the function of the claustrum?* Philosophical Transactions of the Royal Society B 360 (1458): 1271-1279.

Kondepudi, D, Kay, D & Dixon, J. (2017) *Dissipative Structures, Machines, and Organisms: A Perspective*. Chaos, 27.

- Kriegel, U. (2008) *Real Narrow Content*. *Mind & Language* 23 (3): 304–328
- Kriegel, U. (2011) *Subjective Consciousness: A Self Representational Theory*. Oxford University Press.
- Kriegel, U. (2012) *Self-Representationalism and the Explanatory Gap*, in Perry, J & Liu, J. (2012) *Consciousness and the Self: New Essays*. Cambridge University Press.
- Kriegel, U. (2014) *The Sources of Intentionality*. Oxford University Press.
- Kriegel, U. (2009) *Subjective consciousness: A self-representational theory*. Oxford: Oxford University Press
- Landauer, R. (1961) *Irreversibility and Heat Generation in the Computing Process*. *Journal of Research and Development*, 5 (2): 183-191
- Landauer, R. (1961) *Irreversibility and Heat Generation in the Computing Process*. *IBM Journal of Research and Development* 5 (3): 261-269
- Langton, R & Lewis, D. (1998) *Defining 'Intrinsic'*. *Philosophy and Phenomenological Research*, 58: 333–45
- Langton, R. (2006) *Kant's Phenomena: Extrinsic or Relational Properties. A reply to Allais*. *Philosophy and Phenomenological Research*, 73 (1). Accessed from <http://web.mit.edu/langton/www/pubs/Allais.pdf>, 07/05/2021.
- Langtry, B. (1972) *Identity and Spatio-Temporal Continuity*. *Australian Journal of Philosophy*, 50 (2): 184-189
- Laudan, L & Leplin, L. (1991) *Empirical Equivalence and Underdetermination*. *Journal of Philosophy* 88 (9): 449-472
- Leibniz, G. (1714/2014) *The Monadology*. Edinburgh University Press.
- Lewis, D. (1983) *"Extrinsic Properties"*. *Philosophical Studies*, 44: 197–20

- Lewis, D. (1986) *Philosophical Papers Volume II*. Oxford University Press.
- Levinson, J. (1978). *Properties and related entities*. *Philosophy and Phenomenological Research*, 39(1), 1–22.
- Lightner, D. (1997) *Hume on Conceivability and Inconceivability*. *Hume Studies* 23 (1): 113-132
- Lipton, P. (1991/2004) *Inference to the best explanation*. Routledge.
- Loar, B. (1990/2017) *Phenomenal States*, in Loar, B. (2017) *Consciousness and Meaning: Selected Essays*. Oxford University Press, pp. 195-220
- Lorenzini, G, Moretti, S & Conti, A. (2011) *Fin Shape Thermal Optimisation Using Bejan's Constructal Theory*. Morgan & Claypool.
- Lososky, M. (1987) *Individuation and the Bundle Theory*. *Philosophical Studies*, 52 (2): 191-198.
- Martin, C. B. (1993). Power for realists. In J. Bacon, K. Campbell, & L. Reinhardt (Eds.), *Ontology, cause and mind: Essays in honour of D. M. Armstrong*. Cambridge: CUP
- Martin, C. B. (1997). *On the Need for Properties: The Road to Pythagoreanism and Back*. *Synthese*, 112(2), 193-231.
- Martin, C. B. & Heil, J. (1999). *The Ontological Turn*. *Midwest Studies in Philosophy*, XXIII, 34-60.
- Maxwell, J. (1853) 'Reflex Musings', in Campbell, L. (1882) *The Life of James Clerk Maxwell*. Macmillan and Co.
- Maxwell J. (1862) "From a letter to Lewis Campbell", in Harman P. (1990) *The Scientific Letters and Papers of James Clerk Maxwell* (1) 1846-1862: 711-2.
- Maxwell, J. (1862) in Stanley, M. (2015) *The Pointsman: Maxwell's Demon, Victorian Free Will and the Boundaries of Science*. *Journal of the History of Ideas* 69 (3): 467-491

- Maxwell, J. (1867), in Knott, C. (1911) *Life and Scientific Work of Peter Guthrie Tait*.  
Cambridge: Cambridge University Press.
- Maxwell, J. (1871/2001) *The Theory of Heat*. New York: Dover
- McDowell, J. (1996) *Mind and World*. Harvard University Press.
- McGinn, C. (2006) *Hard Questions*, In Strawson, G et al. (2006) *Consciousness and Its Place in Nature: Does Physicalism Entail Panpsychism?* Imprint Academic, pp. 90–99
- Michaelian, K. (2017) *Microscopic Dissipative Structuring at the Origin of Life*. *Heliyon*, 3 (10).
- Miller, G. (2018) *Can subjects be proper parts of other subjects? The de-combination problem*.  
*Ratio* 31 (2): 137-154
- Molnar, G. (2003/7) *Powers: A Study in Metaphysics*. Oxford University Press.
- Montague, M. (2014) *Evaluative Phenomenology*, in Roeser, S & Todd, C. (2014) *Emotion and Value*. Oxford University Press.
- Montero, B. (2016) *What Combination Problem?*, in Bruntrup, G & Jaskolla, L. (2016) *Panpsychism, Contemporary Perspectives*. Oxford University Press.
- Moore, D. (2012) *On Robinson's Response to the Self-Stultifying Objection*. Review of  
*Philosophy and Psychology* 3 (4): 627-641.
- Moore, G. (1903) *The Refutation of Idealism*. *Mind* 12 (48): 433-453
- Morch, H. (2014) *Panpsychism and Causation: A new argument and a solution to the combination problem*. Unpublished PHD thesis, University of Oslo. Accessed from  
<https://www.newdualism.org/papers/H.Morch/Morch-dissertation-Oslo2014.pdf> on  
08/06/2020.
- Morch, H. (2018) *Does Dispositionalism entail Panpsychism?* *Topoi*:1-16
- Moreno, E & Rhiner, C. (2014) *Darwin's multicellularity: from neurotrophic theories and cell competition to fitness fingerprints*. *Current Opinion in Cell Biology* 31: 16-22.
- Nagasawa, Y. (2010) The Knowledge argument and epiphenomenalism. *Erkenn*, 72: 37-56

- Nagasawa, Y & Wager, J. (2016) *Panpsychism and Priority Cosmopsychism*, in Bruntrup, G & Jaskolla, L. (2016) *Panpsychism: Contemporary Perspectives*. Oxford University Press.
- Nagasawa, Y & Wager, K. (2016) *Panpsychism and Priority Cosmopsychism*, in Bruntrup, G & Jaskolla, L. (2016) *Panpsychism Contemporary Perspectives*. Oxford University Press.
- Nagel, T (1974) *What is it like to be a Bat?* *The Philosophical Review* 83 (4): 435–450.
- Nagel, T. (1979) *Mortal Questions*. Cambridge University Press.
- Nagel, T. (2012) *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature is Almost Certainly False*. OUP USA.
- Nani, A; Manuello, J; Mancuso, L; Liloia, D; Costa, T; Cauda, F. (2019) *The Neural Correlates of Consciousness and Attention: Two Sister processes of the Brain*. *Frontiers in Neuroscience* 13 (1169).
- Nesic, J. (2018) *Subjectivity Emergence in Panpsychism*. *Facta Universitatis* 18 (3): 169-178
- Niiniluoto, I. (2003) *Critical Scientific Realism*. Oxford University Press.
- Olson, R. (1987) *An essay on facts*. Stanford: CSLI.
- Papinaeu, D. (2016) *Panpsychist Infusion*, in Bruntrup, G & Jaskolla, L. (2016) *Panpsychism: Contemporary Perspectives*. Oxford University Press
- Papineau, D. (2001) *"The Rise of Physicalism"*, in Loewer, B. (2001) *Physicalism and its Discontents*. Cambridge University Press.
- Papineau, D. (2002) *Thinking About Consciousness*. Oxford University Press.
- Paul, L. (2012) *Metaphysics as Modeling: The handmaiden's tale*. *Philosophical Studies* 160: 1-29
- Peat, D. (1987) *Synchronicity: The Bridge between Matter and Mind*. Bantam Doubleday Dell Publishing Group.
- Peirce, C. (1857-1866/1982) *Writings of Charles S. Peirce: Chronological Edition, Volume 1: 1857-1866*. Indiana University Press.

- Peirce, C. (1877/2014) *Illustrations of the Logic of Science*. Open Court Publishing Co.
- Pelczar, M. (2014) *Physical Time, Phenomenal Time, and the Symmetry of Nature*, in Oaklander, L. (2014) *Debates in the Metaphysics of Time*. Bloomsbury.
- Pernu, T. (2017) The Marks of the Mental. *Frontiers in Psychology*, 8 (1084).
- Pigliucci, M & Goff, P. (2020) *On the Problem of Consciousness and Panpsychism*. Accessed from [https://letter.wiki/conversation/277#letter\\_1471](https://letter.wiki/conversation/277#letter_1471) on 03/03/2020.
- Pinker, S. (2019) *Enlightenment Now: The case for Reason, Science, Humanism and Progress*. Penguin.
- Pitt, D. (2004) *The Phenomenology of Cognition, Or, What is it like to think that P?* *Philosophy and Phenomenological Research*, 69 (1): 1-36
- Prigogine, I. (1993) *Order Out of Chaos: Man's New Dialogue with Nature*. Flamingo.
- Popovic, M. (2017) *Living organisms from Prigogine's Perspective*. *Journal of Biological Education* 52 (3).
- Reardon, S. (2017) *A Giant Neuron Found Wrapped around Entire Mouse Brain*. *Nature* 543: 14-15
- Robb, D., & Heil, J. (2014). *Mental Causation*. In Zalta, E. (2014) *Stanford Encyclopedia of Philosophy*. Spring 2014 Edition
- Robinson, W. (1982) "*Causation, Sensations and Knowledge*". *Mind* 91: 524–540
- Robinson, W. (2006) "*Knowing Epiphenomena*". *The Journal of Consciousness Studies* 13: 85–100
- Robinson, W. (2013) "*Experiencing is Not Observing: A Response to Dwayne Moore on Epiphenomenalism and Self-Stultification*". *The Review of Philosophy and Psychology* 4(2): 185–19
- Robinson, W. (2015) *Russellian Monism and Epiphenomenalism*. *Pacific Philosophical Quarterly* 99 (1): 100-117



- Robinson, W. (2018) *Epiphenomenal Mind: An Integrated Outlook on Sensations, Beliefs, and Pleasure*. Routledge
- Rosenberg, G. (2004) *A Place for Consciousness: Probing the deep structure of the natural world*. Oxford University Press
- Rosenthal, D. (1991). *The independence of consciousness and sensory quality*. Philosophical Issues 1: 15–36.
- Rosenthal, D. (2001) *Introspection and self-interpretation*. Philosophical Topics 28(2): 201–233
- Rosenthal, D. (2005) *Consciousness and Mind*. Oxford University Press.
- Russell, B. (1912/1997) *The Problems of Philosophy*. Oxford University Press
- Russell, B. (1948/2009) *Human Knowledge: Its Scope and Limits*. Routledge.
- Schaffer, J. (2010) Monism: The Priority of the Whole. Philosophical Review 119 (1): 31-76.
- Schneider, S. (2017) What Breathes Fire into the Equations? Journal of Consciousness Studies 24 (9-10).
- Schrodinger, E. (1944/2012) *What is life? With Mind and Matter and Autobiographical Sketches*. Cambridge University Press.
- Schwartenbeck, P; Fitzgerald, T; Dolan, R & Friston, K. (2013) *Exploration, novelty, surprise, and free energy minimisation*. Frontiers in Psychology 4 (710).
- Seager, W. (2006) *The 'Intrinsic Nature' Argument for Panpsychism*. Journal of Consciousness Studies 13 (10-11).
- Seager, W. (2010) *Panpsychism, aggregation and combinatorial infusion*. Mind and Matter 8 (2): 167-184
- Searle, J. (2004) *Mind: A Brief Introduction*. Oxford University Press.
- Segal, A & Goldschmidt, T. (2017) The Necessity of Idealism, in Goldschmidt, T & Pearce, K. (2017) *Idealism: New Essays in Metaphysics*. Oxford University Press.

- Shani, I. (2015) *Cosmopsychism: A Holistic Approach to the Metaphysics of Experience*.  
Philosophical Papers (44) 3: 389-437.
- Shani, I. (2015) 'Cosmopsychism: A Holistic Approach to the Metaphysics of Experience'.  
Philosophical Papers 44(3): 389–437
- Shannon, C. (1948) *A Mathematical Theory of Communication*. The Bell System Journal 27:  
379-434
- Sider, T. (2009) 'Ontological realism,' in Chalmers, D, Manley, D & R. Wasserman. (2009)  
*Metametaphysics*, Oxford University Press, 384-42.
- Sider, T. (2012) *Writing the Book of the World*. Oxford University Press.
- Sider, T. (2014) *Consequences of Collapse*, in Baxter, D & Cotnoir, A. (2014) *Composition as  
Identity*. Oxford University Press.
- Siewert, C. (1998) *The Significance of Consciousness*. Princeton University  
Press.
- Skow, B. (2007) *Are Shapes Intrinsic?* Philosophical Studies (133):111-130
- Skrbina, D. (2007) *Panpsychism in the West*. MIT Press
- Smith, E. (2008) Thermodynamics of Natural Selection I: Energy flow and the limits of  
organization. *Journal of Theoretical Biology* 252: 185-197
- Smuts, J. (1926/2013) *Holism and Evolution*. Gestalt Journal Press
- Spencer-Brown, L. (1969) *Laws of Form*. Allen and Unwin.
- Stanley, M. (2015) *The Pointsman: Maxwell's Demon, Victorian Free Will and the Boundaries  
of Science*. *Journal of the History of Ideas* 69 (3): 467-491.
- Stoljar, D. (2010) *Physicalism*. Routledge.
- Strawson, G. (1994) *The Mental Reality*. MIT Press.

- Strawson, G. (2006) *Why Physicalism entails Panpsychism*, in Strawson, G. (2006) *Consciousness and its place in nature: Does Physicalism entail Panpsychism?* Imprint Academic, pp. 3-30.
- Strawson, G. (2008) *Raw Materialism: And other Essays*. Oxford University Press.
- Strawson, G. (2011) *The Minimal Subject*, in Gallagher, S. (2011) *The Oxford Handbook of the Self*. Oxford University Press.
- Strawson, G. (2012) *Real Naturalism*. Proceedings and Addresses of the American Philosophical Association 86 (02): 125-154
- Strawson, G. (2018) *The Subject of Experience*. The Philosophical Quarterly 68 (273): 870–873.
- Szilárd, L. (1929) *On the Decrease in Entropy in a Thermodynamic System by the Intervention of Intelligent Beings*, in Harvey, S. & Rex, A. (2014) *Maxwell's Demon: Entropy, Information, Computing*. Princeton University Press, pp. 124-34.
- Szilard, L. (1929) *On the Reduction of Entropy in Thermodynamic System by the Intervention of Intelligent Beings*. Zeitschrift für Physik 53 (11-12): 840-856
- Tartaglia, J. (2008) Intentionality, Consciousness, and the Mark of the Mental: Rorty's Challenge. The Monist, 91 (2), 324-346
- Tahko, T. (2014) Boring Infinite Descent. Metaphilosophy 45 (2): 257-269
- Taylor, H. (2018). *Powerful qualities and pure powers*. *Philosophical Studies*, 175(6), 1423–1440.
- Taylor, H. (2019) *Powerful Problems for Powerful Qualities*. *Erkenn*. Accessed from <https://doi.org/10.1007/s10670-019-00199-y> (08/09/2020).
- Teller, P. (1986) “*Relational Holism and Quantum Mechanics*”. *British Journal for the Philosophy of Science*, 37: 71–81
- Tononi, G. (2004) *An information integration theory of consciousness*. *BMC Neuroscience* 5 (42).

- Torday, J. (2015) Homeostasis as the Mechanism of Evolution. *Biology* 4 (3): 573-590
- Torday, J & Rehan, V. (2017) *Evolution, the Logic of Biology*. Wiley
- Tye, M (2000) *Consciousness, Color, and Content*. MIT Press
- Walker, S & Davies, P. (2016) *The 'Hard Problem' of Life*, in Walker, S. Davies, P and Ellis, G. (2016) *From Matter to Life: Information and Causality*. Cambridge University Press.
- Walker, S. (2017) *Origins of life: problems in physics, a key issues review*. Reports on progress in physics 80 (9).
- Weatherson, B. (2001) *Intrinsic Properties and Combinatorial Principles*. *Philosophy and Phenomenological Research*, 63: 365–8
- Williford, K. (2015). *Representationalisms, Subjective Character, and Self-Acquaintance*, In Metzinger, T & Windt, J (2015) *Open MIND*: 39(T). Frankfurt am Main: MIND Group
- Wolfram, S. (2002) *A New Kind of Science*. Wolfram Media Inc.
- Zahavi, D and Kriegel, U. (2015) *For-me-ness: What it is and what it is not*, in Dahlstrom, D, Elpidorou, A and Hopp, W. (2015) *Philosophy of Mind and Phenomenology: Conceptual and Empirical Approaches*. Routledge: London.
- Zahavi, D. (2005) *Subjectivity and selfhood: Investigating the first-person perspective*. Cambridge: MIT Press.