

The One and the Many of Experience

*The tension between mereology and holism
in Tim Bayne's account of the unity of consciousness.*

Antonios Kaldas
MB,BS(Syd), GradDipPhil(Mac)

Department of Philosophy
Macquarie University, Sydney, Australia

A thesis submitted in fulfilment of the requirements for the degree of
Masters of Research
Submitted 9 October 2015

Contents

Abstract	5
Declaration	7
Acknowledgments	9
Introduction	11
Reflections on methodology	13
Some notes on terminology	17
Chapter One. Tim Bayne's account of phenomenal unity	19
1.1 Phenomenal unity thesis.....	19
1.2 Bayne's mereology	23
1.3 Bayne's holism	26
1.4 Previous objections to the mereological conception.....	27
1.5 Summary	40
Chapter Two. The tension between mereology and holism	43
2.1 Bayne's holism revisited	45
2.2 Bayne's mereology revisited	46
2.3 Bayne's "parts"	53
2.4 Can the tension be resolved?	60
2.5 Summary	72
Chapter Three. Phenomenal interdependence	75
3.1 Phenomenal interdependence and its 'strengths'	76
3.2 Bayne and phenomenal interdependence	87
3.3 On the grounding question	89
3.4 On the parts question	98
3.5 Summary	100
Conclusion	103
Bibliography	107

Abstract

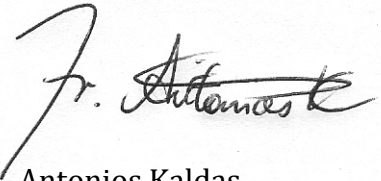
When eating a grape while listening to a symphony, you have a single phenomenally unified experience containing both component experiences, rather than many disjoint ones. Tim Bayne's account of the unity of consciousness has two features that seem to be in tension with each other. On the one hand, he insists that component experiences are real token phenomenal experiences in their own right (mereology) and that phenomenal unity is grounded in mereological relations. On the other hand, he insists that the whole experience is metaphysically prior to the components and that component experiences cannot exist independently of the whole (holism). After briefly surveying previous objections to Bayne's mereological conception, I elaborate on this tension by comparing Bayne's holism with the "No Experiential Parts" view, and consider a number of interpretations of Bayne that might resolve the tension. Perhaps it is just a matter of different descriptions? Perhaps there is a special kind of decomposition? Or perhaps the mutual influence of component states upon each other's characters can support Bayne's account? Ultimately, none of them seem satisfactory, and the tension remains.

Declaration

I certify that the work in this thesis has not been submitted for a degree nor has it been submitted as part of requirement for a degree except as fully acknowledged within the text.

I also certify that the thesis has been written by me. Any help that I have received in my research work and the presentation of the thesis itself has been acknowledged.

I certify that all information sources and literature used are indicated in the thesis.



Antonios Kaldas

Acknowledgements

I would like to acknowledge the longsuffering and wise guidance of my supervisor, Dr Rachael Brown, who has modelled the ideal of an academic philosopher as well as anyone I have met, and who has so capably moulded my philosophical skills from somewhat unpromising clay. I have also benefitted significantly from the advice and support of Associate Professor Richard Menary who started this road with me and Dr Colin Klein with whom I look forward to continuing it, as well as from discussions in the *Macquarie University Consciousness Reading Group*, run by Dr Glenn Carruthers. I would like to thank Professor Jeanette Kennett for her consistent support and encouragement throughout my journey in academic philosophy, and Professor Nicolas Smith the department's HDR Supervisor for his assistance in my moments of need.

My family have not only given me time off my domestic responsibilities to work on this thesis, but have also consistently challenged my ideas and inspired new ones. I am deeply grateful to Dalia, Samuel and Bethany.

*“When I consider the mind, that is to say,
myself inasmuch as I am only a thinking thing,
I cannot distinguish in myself any parts,
but apprehend myself to be clearly one and entire”*

(Descartes, 1641, p. 196).

*“Granting the oneness to exist, what facts will be different
in consequence? What will the unity be known-as?”*

(James, 1910).

Lecture IV: *The One and the Many*.

Introduction

You see a blue circle. You have two visual experiences that seem clearly different experiences: an experience of blueness, *and another* of circularity. And yet, it seems impossible to pull these experiences apart—they form a single, unified and integrated experience. Why is this so? This thesis is an investigation of synchronic phenomenal consciousness, the experience we have at any given moment in time, and about the phenomenal unity of this consciousness—our apparently universal experience that at any moment in time we seem capable of having only a single unified experience rather than a divided experience.

Tim Bayne’s fascinating and fecund account of phenomenal unity in his book, *The Unity of Consciousness* (2010) and elsewhere (see below) moves our understanding forward greatly, but at its heart, there seems to be a serious contradiction. Bayne considers phenomenal consciousness at a moment in time to have genuine parts (mereology) and grounds its unity in a relation of subsumption between the parts and the whole. But contrary to how subsumption relations are normally thought

of, he also argues that we should think of the parts as deriving their existence from that single whole rather than the other way around (holism). In wanting us to see experience as both one and many, is he asking for too much? This thesis explores this apparent tension between Bayne's mereology and his holism and proposes and evaluates some possible paths to resolution.

I begin in Chapter One by sketching three relevant features of Bayne's account: the *phenomenal unity thesis*; *mereology*; and *holism*. I then consider eight objections to his mereological conception of phenomenal unity, together with his responses to tease out the nuances of his subtle account. The last of these, the ontological profligacy objection, raises the apparently unnecessary extravagance involved in thinking that *both* parts and wholes are token phenomenal experiences. This view flows necessarily from Bayne's decision to adopt a *tripartite conception* of experiences, whereby an experience is individuated using three criteria: the *subject* of the experience; the *time* at which it occurs; and its *phenomenal character*. Accordingly, the parts of synchronic experience on this kind of mereology are capable of existence apart from the whole experience. But this view, I propose, sits in an unresolvable tension with Bayne's holism.

To see this, I compare the paradigmatically holistic *no experiential parts* model (NEP) to Bayne's holism in Chapter 2, and identify two important points of divergence between them: the *parts question* and the *grounding question*. Where NEP is the view that synchronic experience has no parts that are experiences themselves, Bayne's mereology requires that there be such parts, but his holism requires that they be metaphysically dependent on the whole experience, what I

shall call *non-detachable parts*. But non-detachable parts are metaphysically odd and apparently unique to Bayne's view. I use the two questions above to evaluate two possible interpretations of Bayne's account that might make sense of non-detachable parts, but find both wanting.

In Chapter Three I introduce *phenomenal interdependence*, the idea that component experiences influence each other's phenomenal character, and use it to show that NEP views make more sense of the mechanism of phenomenal unity than does Bayne's holism. I also explore whether phenomenal interdependence might be used to make sense of non-detachable parts, but find it cannot. In the absence of a successful interpretation, I conclude that Bayne's mereology and his holism cannot successfully be synthesised into a single coherent model of phenomenal unity.

The Conclusion rounds off the discussion with a summary and some reflections on the broader implications of this critique of Bayne's account for more general questions such as how we ought to individuate experiences, and how a better understanding of the nature of phenomenal unity can constrain models of phenomenal consciousness itself.

Reflections on methodology

Phenomenal consciousness is by its ephemeral nature a difficult concept to define and to investigate. In this thesis the standard methodologies of reasoning and argumentation—conceptual analysis, analogy, introspection, thought experiments,

and intuition—are employed to differing extents as appropriate to the subject matter. The chief hypothesis of this thesis is that there is a tension between two aspects of Bayne’s account of phenomenal unity, so the best way to explore this tension and see if it can be relieved is *conceptual analysis* (Beaney, 2015). Analysing concepts like “experience”, “parts”, “mereology” and “holism” teases out this tension and suggests possible avenues for resolving it. The ephemeral nature of phenomenal conscious makes it difficult to reason about, so *analogies* that are more familiar to us would be very useful (Thagard, 2005). Unfortunately, there seems to be nothing else in the world that is sufficiently similar to consciousness to justify its use as an effective analogy. However, there are many analogies to Bayne’s mereology, some of which I employ below. *Modelling* (Weisberg, 2013) is an important form of analogical thinking that is frequently used in consciousness studies. The models I introduce, however, are relatively simple and informal, intended only to make very specific points relevant to my argument and not intended to be in any way comprehensive.

Introspection deserves special attention, since it is the only direct access we have to phenomenal consciousness. Bayne (2010, p. 95) states that the “received view within consciousness studies is that the only legitimate tool for measuring consciousness is introspective report”. Discussions of the unity of consciousness almost always employ passages describing an introspective episode to set the scene for the analyses that follow (e.g., Bayne, 2010, p. 5; Tye, 2003, pp. 2–5).

But there is ongoing debate as to the reliability of introspection. The respected journals, the *Journal of Consciousness Studies*¹ and *Consciousness and Cognition*,² both dedicated whole issues exclusively to discussing it against the background of how we understand consciousness. Overgaard (2006) provides an excellent and brief overview of these recent discussions. One relatively safe role introspection can play is that of furnishing the raw data to be subjected to analysis or as a kind of intuition pump, a pointer to interesting leads worth exploring more fully using other tools. This is how I employ it below.

Two other ways to generate fodder for philosophical argument feature below. *Intuition* (Pust, 2012) is both important and useful but like introspection must be used with care. There is ongoing debate as to just what intuitions are (e.g., Carruthers, 2015) and there are many ways intuitions can lead one astray. There has been a spectrum of views on the soundness of intuition, ranging from those who see it as epistemologically valueless (Cummins, 1998), to those who argue that intuitions provide at least *prima facie* justification to believe in their contents (Bonjour, 1998). Intuition generally plays the latter role in most of the modern discussions of phenomenal unity I have read, and that is how it is used in this

¹ Volume 10, Numbers 9-10, 2003 and Volume 11, Numbers 7-8, 2004.

² Volume 15, Number 4, December 2006: "Special Issue on Introspection".

³ Famous thought experiments about phenomenal consciousness include the *Chinese Room* (Searle,

² Volume 15, Number 4, December 2006: "Special Issue on Introspection".

paper. *Thought experiments*³ (Gendler, 2007), like Avicenna's *Flying Man* which I adapt below, allow us to explore situations in thought that are impractical or impossible to explore in reality. But of course, this benefit also has a downside. There may be good reasons why the fictional situation is impossible in reality, reasons that invalidate any results of the thought experiment from being applied back to reality. Indeed, positions on the validity of thought experiments range from outright condemnation of Duhem (Brown & Fehige, 2014), through deep suspicion and insistence on adhering to scientific possibility (Wilkes, 1988), normative criteria for thought experiments (Häggqvist, 2009), to positive endorsement of their use in personal identity cases (Hershenov, 2008). Nonetheless, even the strongest endorsements fall short of claiming that any thought experiment can comprehensively establish the truth of a position on consciousness, so I employ them, like intuitions, more as fodder for arguments rather than as proofs.

Finally, much philosophy of mind today engages empirical data quite deeply, something at which Bayne excels. In this thesis however, I have reluctantly focused on the conceptual coherence of Bayne's view and left the empirical dimension for another time.

³ Famous thought experiments about phenomenal consciousness include the *Chinese Room* (Searle, 1980), *Mary the Scientist* (Jackson, 1982) and what it might be like to be an echolocating bat (Nagel, 1974).

Some notes on terminology

I use the terms “unity of consciousness”, “phenomenal unity” and occasionally “co-consciousness” generically for the purposes of this thesis, simply referring to the concept of unified phenomenal consciousness without in any way suggesting a preference for the different specific theories of phenomenal unity that have been proposed under those names (Brook & Rayment, 2014). Similarly, I mostly use the term “phenomenality” to denote “phenomenal consciousness”, what Nagel (1974) famously characterised as “what it is like” to have an experience. I use the terms “experience”, “phenomenal state” and “phenomenal mental state” interchangeably. “Parts”, “components” and “component experiences” are likewise interchangeable. Unless otherwise indicated in the text, I use these terms in a general manner without pre-judging or committing to any one view about their nature. Thus, when I use the word “part” I do not mean thereby to imply support for the Experiential Parts view of phenomenal unity, or for Bayne’s theory of the parts of experience (see below), but merely to indicate an entity whose true nature needs to be further defined, which I endeavour to do in the text.

Unless otherwise stated I am always referring to phenomenal experience at a moment in time (synchronic) rather than through time (diachronic). The duration (or lack thereof) of synchronic experience does not affect my arguments. For the sake of clarity I have used a common notation of my own throughout for most formal discussions of states and components, rather than the manifold notations of the original authors. Generally, a capital letter (e.g. “*A*”, “*M*”) denotes a token phenomenal experience. Since this thesis is chiefly about Bayne’s account in his

book, *The Unity of Consciousness* (2010), wherever a page or chapter reference appears without citation of the author or year, it refers to that book.

1. Tim Bayne's account of phenomenal unity

In this chapter I sketch the features of Bayne's account of phenomenal unity that are relevant to my argument: the Phenomenal Unity Thesis; mereology and subsumption; and holism. I then summarise the chief objections to his account raised so far in the literature that are relevant to my argument, together with Bayne's responses. This lays the foundation for the arguments in Chapters Two and Three that are the principal subject matter of this thesis.

1.1. Phenomenal unity thesis

The question of *phenomenal unity* has come (back) to philosophical prominence in the past two decades.⁴ From the outset, Bayne distinguishes phenomenal unity from other kinds of cognitive unity (pp. 9-11). It is not, for example, *subject unity*,

⁴ For the history of thought on the unity of consciousness since Early Modern times, see Bayne & Chalmers (2003), Brook & Raymond (2014), and Zahavi (2011). For ancient sources, see Lorenz (2009).

which is the possibly trivial thesis that anything you experience is unified by virtue of its being experienced by a single subject. Nor is it *representational* unity, which deals with whether the contents represented in phenomenal experience cohere and connect to each other. Phenomenal unity is specifically about unity on the phenomenal level, rather than the level of the subject or the representational content of experience.⁵

Bayne developed his account of phenomenal unity largely as a means to better understand and to support the *Phenomenal Unity Thesis* (PUT).

“Phenomenal Unity Thesis:

Necessarily, any set of phenomenal states of a subject at a time is phenomenally unified” (Bayne & Chalmers, 2003, p. 33).

The PUT has five elements which I list below. Some are more relevant to my thesis than others.

1. “*Necessarily*”

Bayne argues for a biological necessity rather than a logical necessity (p. 17). That is, phenomenal unity seems necessary in organisms with brains like ours, but it is conceivable that a completely different kind of organism/system could have disunified consciousness. In exploring the nature of phenomenal unity in this thesis, I will of course also be touching on why it should be necessary, and a certain

⁵ See also Bayne and Chalmers (2003) and Tye (2003) for additional discussion.

biological necessity does seem to follow from the account I give below.

2. *“phenomenal states”*

What is the structure of a phenomenal state and how should we individuate phenomenal states? Bayne prefers a tripartite approach whereby phenomenal states are individuated according to three parameters: subject, time and phenomenal character (p. 24). This is a question that is deeply relevant to this thesis.

3. *“a subject”*

The question of how to define the subject of experience is an open one for Bayne (and others) and fraught with difficulties. As a working definition, he defines a subject as an organism for most of his book (p. 16), but he points out the fatal flaws in this definition in Chapter Twelve and proposes instead a “virtual centre of phenomenal gravity” whereby “streams of consciousness are constructed ‘around’ a single intentional object” (p. 289). This question is not central to this thesis, whose arguments apply under most definitions of a subject of experience.

4. *“at a time”*

How ‘long’ is synchronic phenomenal experience, experience “at a time”? Bayne acknowledges the difficulties in defining this, but considers these difficulties not to encroach on the evaluation of the

PUT, nor on his account generally (pp. 17-18). I make the same assumption for the purposes of this thesis.

5. “*phenomenally unified*”

What do we mean by “unified”? What is the nature of this unity? Is it mereological or simple? What exactly is unified, the contents of experience or the act of experiencing itself? That is the major topic in Bayne’s account and of subject of the rest of this thesis.

The seminal 2003 paper, *What is the Unity of Consciousness*, co-authored with David Chalmers, comprehensively developed the concept of *subsumption* as the relation that best describes the unity of phenomenal states, a concept Bayne developed into his *mereological conception* of phenomenal unity in his 2010 book, *The Unity of Consciousness*. Here, he also argued strongly for a holistic view on which explanations of phenomenal unity must begin with the whole experience rather than with the parts. His account has been singularly fertile in stirring discussion.⁶ While the majority of this thesis is centred on Bayne’s book, it also draws on some of his other publications to help elucidate and elaborate his views on certain significant issues.

⁶ Bayne himself has pursued the issue further in various articles, both before the 2003 co-authored piece with Chalmers (Bayne, 2000, 2001; Brook, 2012) and after (Bayne, 2004, 2005, 2007, 2008). Raymont and Brook (2009) feature a discussion of subsumptional as opposed to relational views. For responses to Bayne’s book see Bennett and Hill (2014b) and Dainton (2014), the discussions in the journal *Analysis* between Bayne (2014a, 2014c) and Hill (2014), Masrour (2014b), and van Gulick (2014b), the *Stanford Encyclopedia of Philosophy* article on the “Unity of Consciousness” (Brook & Raymont, 2014), as well as reviews and responses by Baker (2011) and Mendelovici (2013). The chief texts on the topic of the unity of consciousness are listed in Brook and Raymont (2014, sec. 1).

At the heart of the PUT for Bayne is the age-old fascination with *singularity* and *multiplicity*. Whether a thing is one or many, and the relation of the whole to the parts has been a topic of philosophical interest at least as far back as the pre-Socratics, and right up to the contemporary branch of analytical philosophy known as mereology (Varzi, 2015). This tension between singularity and multiplicity applies to our daily experience of phenomenal consciousness. When I have a visual experience of a blue circle, is that a single experience, or is it two experiences: one of blueness and one of circleness? Or might it even be both at the same time? Bayne's mereological conception can be seen as an attempt to make sense of this ambiguity, and in so doing, ground the necessity clause of the PUT. How does he do this? For convenience, I summarise his approach in two steps: his mereology and his holism.

1.2. Bayne's mereology

Is there a principled way by which we can tell one experience from another? Bayne argues for a *tripartite conception* for individuating experiences (Chapter Two). This means that just three parameters individuate experiences: *subject*, *time*, and *phenomenal character* (p. 24).⁷ On this basis, my synchronic experiences of the taste of a grape and of the sound of a Beethoven symphony, while identical in

⁷ Dainton also discusses this approach but prefers replacing *subject* with "*physical basis*" (2000, p. 186) while Mendelovici (2013) challenges the validity of the approach, especially with regard to how Bayne uses it to discount two-stream models of divided consciousness.

subject and time, have different phenomenal characters and should therefore be counted as two discrete token phenomenal experiences. This approach results in a multiplicity of token experiences within a subject's unified phenomenal experience, and thus opens the door for an approach that deals with parthood—mereology. On Bayne's account then, phenomenal unity is grounded in a relationship between token experiences that comprise the total phenomenal experience (p. 21). The phenomenal state of a single subject at a moment in time that subsumes all the simultaneous individual phenomenal states that subject is experiencing is itself a token phenomenal state, and may be called the subject's *total conscious state* (p. 26), *phenomenal field* (Bayne & Chalmers, 2003, p. 27), *total phenomenal state* (Bayne & Chalmers, 2003, p. 32), *total experience* (Dainton, 2000, p. 95), or *maximal state* (Tye, 2003, p. 40).⁸ What all of these terms share in common is the idea of a token phenomenal state that is not a component of any other phenomenal state except itself. I will be consider these terms to be interchangeable for the purposes of this thesis, and for brevity, will sometimes refer to this maximal state simply as “the whole” and to its component states as its “parts”. Importantly, the total phenomenal state is, like the component states, a token phenomenal state, according both to Bayne and others' reading of Bayne.

“My current conscious state is an experience in its own right, as
are the various fine-grained conscious states—the pain in my left

⁸ Tye actually speaks of a “*maximal PANIC state*”, “states having a poised, abstract, nonconceptual content”. It is the “maximal” component of his terminology that is of interest here, representing the same idea as the other terms: an experiential state that is not a component of any other experiential state.

leg; my olfactory experience of the coffee ... —that are contained within it” (p. 23).

“the conjunction of all of a subject’s phenomenal states at a given time is itself a phenomenal state” (Peacocke, 2013, p. 328).

“A total conscious state is not just a conjunctive state consisting of the subject’s individual phenomenal states. It is a distinct state of phenomenal consciousness, with its own phenomenal character” (Alter, 2010, pp. 27–28).

Bayne proposes that the relation that best fits this relationship between the parts and the whole of phenomenal experience is the relation of *subsumption* (Chapter Two). Subsumption is a relation by which one or more entities are “included” or “absorbed” in something else. An entity at a higher level subsumes one or more entities at a lower level. Thus phenomenal subsumption consists in one phenomenal state subsuming one or more other phenomenal states. The subsumed states thus become components of the subsuming state, although not, as we shall see later, in the way that bricks become components of a wall—there is a subtlety to phenomenal subsumption that requires us to tread carefully in interpreting it. An example is helpful here and I will return to it often.

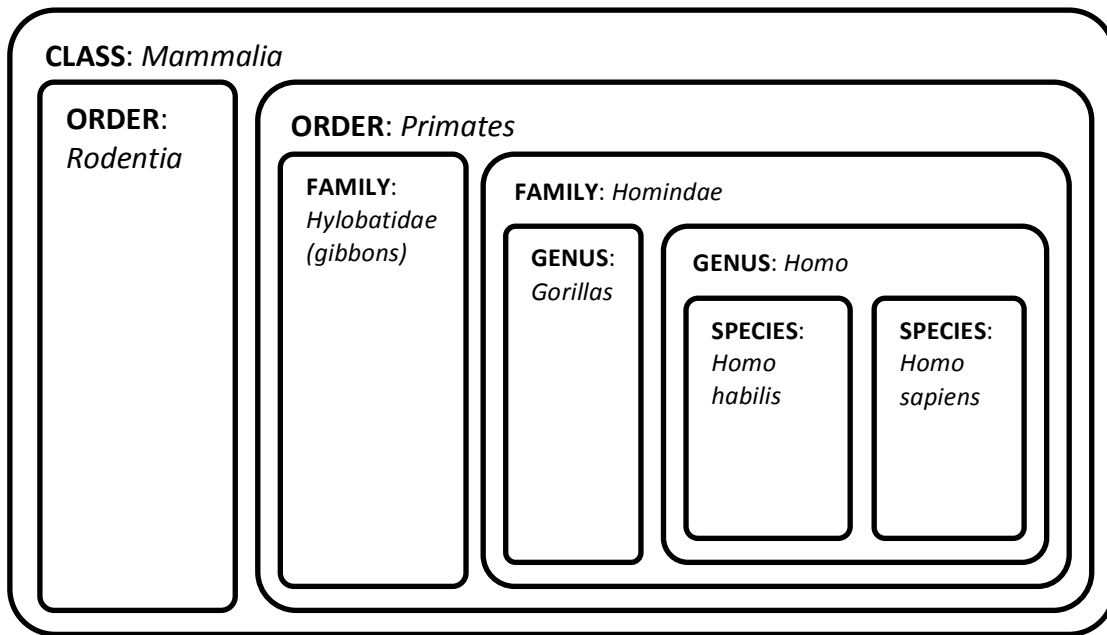


Figure 1 Partial Taxonomy of Life. Adapted from Wilson and Reeder (2005).

The taxonomy of life is a classic case of subsumption. In the taxonomy of living creatures (Figure 1), the *class* “Mammalia” subsumes the *orders* “rodentia” and “primates” among others (not shown). “Primates” in turn subsumes the *families* “hylobatidae” (gibbons) and “hominidae” among others, and so on down the taxonomic tree. Expressed in the opposite direction, we may say that the species “homo habilis” and “homo sapiens” are both subsumed by the genus “homo”, and so on up the taxonomic tree.

1.3. Bayne’s holism

While Bayne insists on our total phenomenal state being composed of components that are themselves token phenomenal states, in Chapter Ten he nonetheless

rejects *atomistic* “building block” models (Searle, 2000), on which the whole derives from the combination of the parts. Instead he prefers a *holistic* view on which the parts exist *only* as parts of the whole. In very rough terms, *atomism* is the position that token total phenomenal states derive their existence from the existence of component phenomenal states (components are “metaphysically prior” to total states), and that component phenomenal states can exist independently of total phenomenal states. *Holism* is the position that total phenomenal states are metaphysically prior to component phenomenal states, and that component phenomenal states as token phenomenal states cannot exist independently of total phenomenal states. I shall describe his holism in more detail below.

1.4. Previous objections to the mereological conception

In this thesis I focus chiefly on the original part of Bayne’s account, his mereological conception, and how it relates to his tripartite conception and holism. There are a number of possible objections that have been or might be raised against Bayne’s mereological conception. Here I leave aside more general theoretical objections such as eliminativist or deflationary denials of the unified conscious self altogether (e.g. Dennett, 1991; Metzinger, 2011), and empirical challenges such as putative cases of disunified conscious subjects, on both of which I will just quote recent assessments and move on.

“Does consciousness exist? There now seems to be a wide consensus that the answer to this question is, Yes” (Brook & Raymont, 2014, sec. 3)

“We follow the mainstream in assuming that deflationary views about phenomenal consciousness, to the effect that there are no such things as qualia or phenomenal properties or qualitative aspects of experience (e.g., Dennett, 1991), or that phenomenal consciousness is not a proper subject of scientific investigation (Irvine, 2013), miss something” (Robinson, Maley, & Piccinini, 2015, p. 366).

“Are there presumptive counter-examples to the unity thesis? Spelled out as we have spelled it out, we do not know of any” (Brook & Raymont, 2014, sec. 3.2).⁹

Rather, my focus in this thesis is on objections and challenges that have been raised specifically against Bayne’s mereological conception and the concept of subsumption as the defining relation underlying phenomenal unity. I consider eight key objections: (1) the self-undermining objection; (2) the phenomenal bloat objection; (3) the transparency objection; (4) the Jamesian objection; (5) the building block objection; (6) the “too exclusive” objection; (7) the insufficiency objection; and (8) the ontological profligacy objection. This list is not exhaustive, and I have focused on objections relevant to my argument. Bayne himself

⁹ See also (pp. 236-237).

addresses the first five of these objections in Chapter Two; the others are drawn from other sources. Below I briefly consider each of these objections in order to illuminate the subtleties of Bayne's view and enlarge the foundation for the arguments that follow.

1.4.1. Self-undermining objection

Tye (2003) offers what Bayne (p. 29) calls the *self-undermining objection*. The mereological conception assumes that there is a maximal phenomenal state that subsumes all component states, but is itself not subsumed by any other state. Call this maximal state *M*, and its components *A* and *B*. But *A* and *M* must also be unified, and according to the mereological conception, says Tye, that must occur through another state, call it *N*, subsuming *A* and *M*. But now, Tye points out, we have contradicted ourselves: we began by stipulating that *M* cannot be subsumed by any other state, yet we are forced to admit that it is in fact subsumed by *N*. Either the notion of a maximal phenomenal state is incoherent (which all reject) or phenomenal unity cannot be a subsumption relation between experiences.

Bayne parries this objection by ascribing the property of *reflexivity* to subsuming states. That is, the maximal state *M* subsumes itself, thus avoiding the need for *N* altogether. One might question whether it is coherent to speak of *M* being a part of itself, but there seems little reason to deny this. By way of analogy with set theory,¹⁰ it is commonly accepted that “every set is (trivially) a subset of itself, but

¹⁰ For a discussion of the relationship between modern mereology and set theory see Varzi (2015). In brief, mereology does not make the ontological commitments to the reality of parts and wholes

no set is a proper subset of itself" (Smith, 2012, p. 441). If we take talk about parthood and inclusion to be similar to talk about sets, then there is no reason why a total phenomenal state cannot subsume itself, that is, cannot be a part of itself (being a "proper part" of itself is just a definitional matter within set theory and need not trouble us here).

1.4.2. Phenomenal bloat objection

Tye (2003) and Hurley (1998) offer a related objection, which Bayne calls the *phenomenal bloat objection* (pp. 30-32).¹¹ If *A* and *B* are phenomenally unified, then the subject will also have a phenomenal experience of the relation of unity between *A* and *B*. Call this experience *R*₁. But *R*₁ has its own phenomenal character and is presumably unified with *A* and *B*. There must therefore be another phenomenal experience of the unity relation between *A*, *B* and *R*₁: call this *R*₂. This line of thinking draws us into an infinite regress of phenomenal states, which is undesirable to say the least. Bayne's response is that experiencing unified experiences need not include an experience of the actual unity relation itself. According to Bayne's approach, we can experience both the colour and the shape of a blue circle without having the explicit experience: "the colour and shape are unified in my experience". This response may deflect the infinite regress aspect of this objection, but I do not believe it deflects the objection that Bayne's tripartite

that set theory has. This aspect of mereology is what allows Bayne some latitude in how he characterizes his parts of phenomenal experience, and therefore also allows me to explore in this thesis the question of what Bayne's "parts" might actually be.

¹¹ See also related discussions in Bayne (2001) on "the combinatorial objection" responding to Dainton (2000), and in Bayne (2005) on the "explanatory regress" and "phenomenal bloat" objections responding to Tye (2003).

conception entails that there be an additional token phenomenal experience on top of *A* and *B*, namely that of *A* and *B* together—see the ontological profligacy objection below (Section 1.4.8).

1.4.3. Transparency objection

Tye (2003) also offers what Bayne calls the *transparency objection* (pp. 32-34). This is an epistemic objection that points out that introspection gives us access only to the *contents* of experience, but not to the *token experiences* themselves, nor (as above) to the *relations* between token experiences. If phenomenal unity is a relation between token experiences as suggested by Bayne then it should not be something that is accessible via introspection, and therefore we should not experience phenomenal unity.

Bayne challenges the premise that introspection has access only to content, but not to the experiences themselves.

“On the tripartite conception of experience, the contrast between ‘access to experience’ and ‘access to the contents of experience’ is a false one, for in having access to the contents of experience, one also has access to experiences themselves” (p. 33).

For Bayne, introspection can and does give us access to *facts* about experiences, such as *that* one is having a conscious experience of a particular kind. If that is so, why could not introspection also give us access to other facts about our

experiences, including facts about the relations between them? In fact, the contents of experience themselves contain facts about the relations between experiences, so if we concede that introspection provides access to the contents, then it also gives access to facts about relations between token experiences.

1.4.4. Jamesian objection

Raymont and Brook (2009) offer what Bayne (pp. 34-35) calls the '*Jamesian*' *objection*. This may be summarized in the maxim that a "mere combination of experiences is not the experience of a combination" (Raymont & Brook, 2009, p. 575)—just putting experiences together does not produce a phenomenally unified total experience. A line of people standing next to each other thinking of different words does not produce a single unified experience of a sentence, so simple parthood is not enough to produce phenomenal unity.

"It is safer ... to treat the consciousness of the alphabet as a twenty-seventh fact, the substitute and not the sum of the twenty-six simpler consciousnesses" (James, 1909, Lecture V).

Bayne's response to this objection is to agree with the above, but to point out that what this objection refers to is *representational unity*, the coherence and integration of the representational contents of experience. The mereological conception does not refer to representational unity but to *phenomenal unity*.

“We obviously *do* need an account of how the experience of parts generates an experience of the whole of which those parts are parts, and it is plausible to suppose that phenomenal unity will play an important role in such an account, but we should not conflate representational unity with phenomenal unity” (pp. 34-35).¹²

Brook however is not satisfied with Bayne’s response to his and Raymont’s objection.

“Paul Raymont and I [2009] have presented arguments against subsumption and in favour of the Jamesian one-state picture just sketched. Bayne discusses them [34–5] but may not have fully appreciated their force ... Even on the subsumption account, there has to be one experience whose unity does not consist in its being a part of a ‘bigger’ experience. So why posit a bunch of ‘smaller’ experiences, too? To be sure, we have to posit a multiplicity of something—but a multiplicity of objects of experience and ways of being conscious would do everything needed” (Brook, 2012, p. 600).

This suspicion of Bayne’s position is one that I will pursue in this thesis. The deflection of the objection by distinguishing representational unity from phenomenal misses the point that the original objection does not refer to representations, but to the phenomenal experiences themselves: “mere

¹² Italics in the original.

combination of *experiences* is not the *experience* of a combination". What are being referred to here are the phenomenal experiences as such, not just what is represented in them, so Bayne's response seems inadequate. It leaves open an important question: what exactly are the "*meri*" (parts) in Bayne's "mereological conception"? I return to this question in Section 2.4, and in Section 3.3 I question whether Bayne himself is not confusing representational unity with phenomenal unity when I consider the difference between phenomenal integration and phenomenal unity.

1.4.5. Building block objection

Searle (2000) discusses the building block theory of phenomenal unity which forms the basis of what Bayne calls the *building block objection* (pp. 35-36). Building block theories hold that phenomenal unity is the product of cementing together, like bricks in a wall, independently existing experiences. They are paradigmatically *experiential parts* views, but I shall return to this topic in much more detail in Section 2.3. For now, I simply note that Bayne emphatically argues that one can take a mereological approach that is nonetheless *not* a building block approach, and that his own mereological conception is in fact neutral on the question of whether the parts of experience exist independently of the whole (atomism) or not (holism). As we shall see, Bayne does take a stand on this question later in Chapter Ten, on the side of holism. Thus, for him, objections to building block theories have no purchase on his mereological conception, although I will argue below (Section 3.2) that that may not necessarily be the case.

1.4.6. Too exclusive objection

Moving on to further objections not considered in *The Unity of Consciousness*, the sixth objection is what I shall call the “*too exclusive*” objection (Bennett & Hill, 2014a; Hill, 2014). Bennett and Hill’s preferred model of the unity of consciousness is a *unity pluralist account*, which they see as taking a broader view of the overall project than Bayne’s. Rather than Bayne’s single unity-making relation of subsumption lying at the root of all phenomenal unity of conscious experience, they see phenomenal unity as arising out of various unity-making relations like subject unity and representational unity, together with a multiplicity of objects of experience. This approach is in direct opposition to Bayne’s arguments against any of these other unity-making relations being responsible for phenomenal unity (Bayne & Chalmers, 2003).

I suspect that the two sides in this disagreement are really embarked on somewhat different projects. While Bayne is concerned with phenomenal unity specifically, Bennett & Hill are concerned not so much with the unity of *consciousness* as with the wider concept of the unity of *experience*;¹³ a broader term that encompasses both phenomenal and non-phenomenal mental states. Given this broader project, it is unsurprising that they would find Bayne’s subsumption relation inadequate to their task. But this in no way counts against Bayne’s case that on the narrower project of phenomenal unity, subsumption can do all the work that is needed, while the other unity-making relations fail to do so. This analysis is brought out nicely by Tye (2003) who distinguishes phenomenal consciousness (*P-consciousness*) from three other types of consciousness: introspective;

¹³ See Bennett and Hill (2014a) and Hill (2009) for further elaboration on this distinction.

discriminatory and responsive (*I*-, *D*-, and *R*-consciousness respectively). He then goes on to show that while other kinds of unity-making relations—namely object, neurophysiological, subject, spatial, higher-order subject and gestalt unities—may be involved in *I*-, *D*-, and *R*-consciousness, they cannot be responsible for *P*-consciousness. For that we need an account of phenomenal unity.

1.4.7. *Insufficiency objection*

The seventh objection I will call the *insufficiency objection*: basically, the charge that subsumption is not up to the task of fully accounting for phenomenal unity. Hill (2014) is not convinced that Bayne's mereological conception makes any substantive progress to our understanding of phenomenal unity, or that Bayne has shown that the subsumption relation does enough work to entail phenomenal unity.

“To defend his mereological conception of unity and the associated unity thesis, Bayne needs a principled way of distinguishing between logical sums of experiences that genuinely count as unified and logical sums that are merely accidental or syncretic”
(Hill, 2014, p. 505).

He appeals to *Leonard-Goodman fusions* that illustrate the fact that being part of the same set does not necessarily entail being unified. The experiences in my mind and yours belong to a set, but they are not phenomenally unified.

Bayne is cautious in his response to Hill's challenge (Bayne, 2014a). He sees it as raising the more general metaphysical question of how we decide to ascribe components of sets to one set and not another. But he does not delve too deeply into this question himself, concluding that an answer to it lies somewhat beyond the scope of his project in *The Unity of Consciousness*. From this response we may surmise that while Bayne sees the value in pursuing the question, he does not see it as raising any serious challenges to his mereological conception specifically.

Dainton (2014) raises a similar objection. He considers the constraints on co-phenomenality required to make subsumption work properly. We cannot allow unrestricted composition, for then your experiences and my experiences should be just as capable of being phenomenally unified as any two of my own individual experiences. Clearly this does not obtain. But there are only two ways of avoiding so undesirable a result: either we constrain unified experiences as being only and all those that enjoy a "conjoint phenomenology"; or we constrain them as being only and all those that are experienced simultaneously by the same subject. Both these options are problematic. The first runs the risk of relegating subsumption to redundancy: by definition any experiences that enjoy conjoint phenomenality are phenomenally unified. The second is clearly untenable since there is no *a priori* reason for assuming that a subject could not have two simultaneous experiences without them being unified. In fact that is just what Hurley (1998), Schechter (2014) and others have argued.

Bayne could counter Dainton's objection by pointing out that subsumption is not redundant, and he has indeed argued just that (Bayne & Chalmers, 2003; Bayne,

2010). It is true that per the PUT, so far as we know, any experiences that enjoy a conjoint phenomenology are necessarily phenomenally unified, but this still leaves the question open as to *how* they come to be phenomenally unified, or *why* they should be phenomenally unified rather than not. These are substantive questions that justify the introduction of subsumption to answer them.

1.4.8. *Ontological profligacy objection*

The last objection to be considered I will call the *ontological profligacy objection*. This can be seen as a sub-objection of the phenomenal bloat objection above (Section 1.4.2), and it has not, I believe, received the attention it deserves.¹⁴ The core concept of Bayne's subsumptive mereological conception is this idea that we must insist on there being a single, unified, token phenomenal state that subsumes all component states. On the tripartite conception then, since the parts and the whole have different phenomenal characters, we must individuate them as different token phenomenal states. Consider Bayne's summary of the mereological conception.

"The mereological view treats phenomenal unity as a relation between token experiences—that is between particular mental states or events. Some token experiences—such as my headache experience and my auditory experience of the trumpet—are parts of a single composite experience and hence phenomenally unified with each other; other token experiences—such as my experience

¹⁴ And sadly I will not be able to remedy that situation in this thesis.

of the trumpet and *your* experience of it—are not phenomenally unified with each other, for there is no experience that contains both of these experiences as parts” (p. 21).

This passage highlights the difficulty that underlies the ontological profligacy objection: what are the ontological statuses of the subsuming state and the component states? Is the subsuming state something that exists *in addition* to the experiences that compose it? And if so, are we not multiplying existent states unnecessarily?¹⁵ Barry Dainton (2000) raised a similar objection which he called the *combinatorial objection*, and which was critiqued by Bayne.

“So two experiences, e1 and e2, will fuse to form a third experience, e3 ... It seems to follow from this that e1 and e3 will be co-conscious, as will e2 and e3, generating conscious states e4 and e5 respectively. But of course, e5 will be co-conscious with e1–e4, and so on, generating still further states of consciousness ... It seems to follow that any stream of consciousness that has more than one experience at a time will have an indefinite number of experiences at a time ... Dainton himself suggests that the simplest way to solve the problem is to deny that e4 and e5 (and similar constructs) are extra or additional experiences; the resultant state is just e3 under a different name.¹⁶ In other words, the

¹⁵ This objection differs from the phenomenal bloat objection in that the problematic putative additional phenomenal state here is not the relation between the parts, but the whole state itself.

¹⁶ In a footnote in the original article Bayne states that this was a “personal communication”.

combination of an experience with one of its parts just is the experience itself" (Bayne, 2001, p. 84).

In the same book, Dainton takes a roughly building block view of phenomenal unity, so his solution to the problem is to posit that it is only the *parts* of experience that qualify as really existing while the whole is just built out of those parts. This contrasts with Tye's (2003) one-experience view which sees the *whole* as being the truly existing experience while the parts are not true discrete experiences but just multiple contents of experience. Bayne disagrees with both. On the one hand, contra Dainton, he argues that the whole is metaphysically and explanatorily prior to the parts of experience (Chapter Ten). On the other hand, contra Tye, he argues that the parts are genuine token phenomenal experiences and that a mereological approach is therefore justified (Chapter Two). It is this tension between Bayne's mereology and his holism that I will explore in the remainder of this thesis.

1.5. Summary

In this chapter I introduced the features of Bayne's account of phenomenal unity that are salient to my thesis. Starting from the PUT, I described how, based on a tripartite approach to individuating experiences, Bayne takes a mereological approach that sees phenomenal unity as a relation of subsumption between the parts of experience and the whole experience, both of which are genuine token experiences. While he insists that his mereology is neutral on the question of

whether the parts can exist apart from the whole, he argues on other grounds for holism, the idea that the parts exist only as parts of the whole. I then considered eight objections to Bayne's mereological conception together with Bayne's responses, concluding with the ontological profligacy objection that worries about unduly multiplying token phenomenal experiences if we consider both the parts and the whole to be such entities. This raised the question of whether there may not be a serious tension in Bayne's account between his mereology and his holism. In the next chapter I explore this tension more fully.

2. The tension between mereology and holism

In Sections 1.2 and 1.3 I sketched two central aspects of Bayne's views on phenomenal unity, his mereology and his holism. The ontological profligacy objection (Section 1.4.8) highlighted Bayne's insistence that the parts of experience are *token phenomenal experiences themselves*. On the face of it, we can see then that Bayne's mereology and his holism seem to be in tension, a tension I explore in more detail in this chapter. Briefly: on the one hand he describes phenomenal unity in terms of a mereology based on the relation of subsumption between the parts and the whole. Such mereologies usually entail parts that are capable of existing independently of the whole to which they belong. In our taxonomic example, primates could still exist even if they were unrelated to rodentia or mammalia. On the other hand, Bayne argues for a phenomenal holism in which the parts—as token phenomenal experiences—*cannot* exist independently of the whole. There appears on the face of it to be a contradiction between these two aspects of his account: how can component experiences be

token phenomenal experiences in their own right, and yet not be capable of existing in their own right, independently of the whole experience?

Adding to the interest of this question is the contrast between Bayne's mereological conception and one-experience views such as those of William James and more recently, Michael Tye, in which there are *no* parts to be accounted for at all. For the sake of brevity in what follows, I shall use the term *detachable* to signify parts that can exist independently of their whole, and *non-detachable* to signify parts that cannot exist without their whole.

In this chapter I explore this tension in Bayne's account by first considering how Bayne's holism sits within contemporary discussion of phenomenal atomism and holism. I then consider mereology generally, before focusing in on Bayne's mereology of phenomenal unity in particular. Next I introduce another distinction prominent in the contemporary discussion of phenomenal unity that is more clear cut than the confused atomism/holism distinction, that between EP and NEP views (see below), and use it to clarify exactly what Bayne's holism is saying. It emerges that there are two significant points of divergence between the Bayne's holism and NEP, which I apply to the aforementioned tension. I conclude that however we try to resolve this tension, none of the solutions are completely satisfactory.

2.1. Bayne's holism revisited

To see the tension in Bayne's account we need first to delve a little more deeply into what his holism and mereology say. First, his holism. In Section 1.3 I gave very rough definitions of phenomenal atomism and holism. It is difficult to give more than a rough definition because there are in contemporary philosophy a number of accounts of the distinction between atomism and holism with respect to phenomenal unity.¹⁷ Different authors have defined atomism and holism in subtly (and sometimes not so subtly) different ways. For example, Lee (2014a) recently blurred the lines between the two positions by arguing that the anti-holist can reject four central theses of atomism without thereby being forced to accept holism (despite holism and atomism often being described as exhaustive of the options available). The same author notes that "it is fair to say that we lack a detailed account of the different ways this distinction can be drawn and how they are related" (Lee, 2014a, p. 288). Situating Bayne's account of his holism against those of others is a demanding task beyond the scope of this study, especially since Bayne's position on atomism/holism is highly nuanced and difficult to define simply. Further confusion arises from a number of similar and overlapping dichotomies in the literature.¹⁸

¹⁷ It is not possible in this thesis to adequately cover this fascinating topic, but the interested reader is referred to the following non-exhaustive list of discussions: Bayne and Chalmers (2003); Brook and Raymont (Brook & Raymont, 2014; Raymont & Brook, 2009); Dainton (2000, 2014); Gallagher and Zahavi (2012); Koksvik (2014); Lee (2014a, 2014b) and Shoemaker (2003). Searle (2000) uses the term "holistic" sparingly but has much of relevance.

¹⁸ These include the Leibnizian/Newtonian distinction (Bayne, 2014a; Masrour, 2014a, 2014b) and see below (Section 2.3); the bottom-up/top-down distinction (Dainton, 2007, 2014; Hill, 1991) and pp. 20, 226; and the federalist/imperialist structure distinction discussed by Bayne himself on pp. 115-119.

What this amounts to is that there is no simple answer to the question, “is Bayne a holist?” The most appropriate answer would be another question: “which holism?” Thankfully we do not need a comprehensive overview of the different accounts of holism for the purposes of this study, only an accurate picture of Bayne’s particular brand of holism. In Bayne’s summary of how he defines atomism and holism below, the two key concepts of metaphysical priority of the whole and independence of component states are evident. Other nuances of Bayne’s particular account will become clear as the discussion progresses.

“Theorists who adopt an atomistic orientation assume that the phenomenal field is composed of ‘atoms of consciousness’—states that are independently conscious. Holists, by contrast, hold that the components of the phenomenal field are conscious only as the components of that field. Holists deny that there are any independent conscious states that need to be bound together to form a phenomenal field. Holists can allow that the phenomenal field can be formally decomposed into discrete experiences, but they will deny that these elements are independent atoms or units of consciousness” (p. 225).

2.2. Bayne’s mereology revisited

Turning now to Bayne’s mereology, the relevant questions are: what motivates Bayne to adopt a mereological approach? What does this approach entail? What

exactly are Bayne's "parts" of phenomenal experience? And ultimately, where exactly does the problem lie between Bayne's mereology and his holism? I consider the first two questions in this section, the third and fourth questions in Section 2.3. In Section 2.4 I propose and evaluate possible ways that the tension between mereology and holism raised by these questions might be resolved.

First, Bayne's motivations for adopting a mereological approach. In Chapter Two Bayne defends his mereology specifically against Tye's (2003) one-experience view. Since this section is extremely significant to our understanding of the nuance in Bayne's account and the tension in it, it is worth considering in some detail. Tye holds that there is only one token phenomenal experience in a subject at a moment in time—the total phenomenal state—and we should not consider component experiences as token experiences at all. Of course this one experience has multiple contents. Bayne argues that Tye's view is revisionary: our pre-theoretical intuition is to accept multiple experiences being bound together. Therefore we should not contemplate Tye's view unless we have very good reason to (p. 22). Bayne's mereological conception is not revisionary in this way but rather respects the pre-theoretical intuition of multiple experiences. However, I argue that there is indeed good reason to think that the pre-theoretical intuition is wrong. Historically in contemporary philosophy the idea of a discrete phenomenal mental state (like tasting a grape) was very useful in debates over questions where this intuition didn't really matter. In debating the merits of, say, substance dualism, behaviourism, identity theory, functionalism or eliminativism as the model that best captures the nature of consciousness itself, few problems if any arise from using discrete "experiences" as fodder for the debate. For this purpose, what

applies to a total state is likely also to apply to any part of that state. For example, if on an identity theory of consciousness, the mental state of tasting a grape is just token identical to the brain state associated with it, the same can be said of the total mental and total brain states of the subject in the same moment. For the purposes of identity theory, it makes little difference to the question of identity between mental and brain states whether you are talking about a total synchronic experience or isolating just part of the experience. It is only when we begin addressing a question like phenomenal unity that a serious problem arises, and we can no longer take it as given that discrete component experiences obtain. Tye is quite right to go back and revise our pre-theoretical intuition about multiple synchronic experiences, for this cuts to the heart of the very question under consideration: the unity of phenomenal consciousness.

Tye goes on to justify his position by appealing to clouds and statues whose parts are not themselves clouds or statues. By analogy the parts of phenomenal experience are not themselves phenomenal experiences, but must be something else. Bayne emphatically rejects this aspect of Tye's account. For Bayne, the parts of experience are *exactly* the same kind of thing as the whole experience, and therefore there must be more than just "one-experience". Bayne justifies his position with four responses to Tye's argument.

First, he points out that clouds and statues are physical objects, whereas experiences are events (p. 23), and that therefore the analogy may not hold. However, Bayne seems to ignore the examples of events given by Tye, such as the eruption of Mount Vesuvius and a movie. Tye argues that a particular flow of lava

is not itself a volcanic eruption, nor is a scene from a movie itself a movie. While these two events are not phenomenal (so far as we can tell) they are still events in all the important ways that make them appropriate as analogies for our purposes for a mental event such as a phenomenal experience.

Second, Bayne points out that while parts of clouds may not *be* clouds, they still *could* be clouds. Further, there are numerous examples where parts of things are things of the same kind themselves: “Arguments contain as proper parts other arguments, battles contain as proper parts other battles” (pp. 22-23). While Bayne makes a valid point, the real question here is not whether clouds could have clouds as parts, but whether phenomenal consciousness is the kind of thing where that is the case, or whether it is the kind of thing where it is not the case, together with things like eruptions and movies. Perhaps this comes down to how you define a cloud, eruption or movie? Normally, a single scene from a movie would not be called “a movie” if my definition of “movie” includes: “is a complete feature-length story”. But if I define “movie” just as any series of photographic frames shown in rapid succession, then clearly even a single scene from the movie qualifies as a movie itself. The real question seems to come down to how we define “experience”. By adopting the tripartite conception, Bayne makes it inevitable that component experiences be considered genuine experiences. Below (Section 2.4.1) I consider the possibility that the difference between Bayne and Tye is merely a verbal one.

Third, Bayne argues that it is the “received view” (Bayne, 2005, 2014b) that a single total token experience cannot capture phenomenal unity. This view is widely accepted by a majority of researchers, and with some claim to being the

default position until someone comes up with strong reasons to dislodge it. He considers Tye to have tried but failed to do exactly that (Bayne, 2005). It is also the intuitive way to think about phenomenal unity. It just seems obvious from introspection that different phenomenal characters means numerically different token experiences, an intuition formalised in his tripartite approach. Further, the contents of our experience can combine and recombine in many different ways. To be capable of such recombination, the most natural way to think of component experiences is as discrete decomposable and recomposable experiences.

“phenomenal contents can occur in different phenomenal contexts.

For example, one can hear a dog barking while: experiencing an itch in one’s right leg; experiencing an itch in one’s left leg; or having no bodily sensations at all ... The intuitive explanation of why one can experience a dog barking without experiencing an itch and vice versa is that the experience of the dog barking is distinct from the experience of the itch” (Bayne, 2005, p. 499).

Further, he points to the fact that component experiences “possess distinct causal profiles and a certain degree of functional autonomy” (p. 248), which provides sufficient reason to consider them as individual entities in any discussion of the nature of phenomenal consciousness. But the question here is going to be how he is to respect that intuition, yet avoid his view becoming a building block view. What kind of decomposition of the whole is Bayne advocating? Below (Section 2.4.2) I consider the possibilities open to Bayne on this question.

Fourth, Bayne argues that because we often find a need to distinguish the parts of a total phenomenal experience—e.g. the auditory part from the gustatory part—there are indeed discreet experiences within a single stream of consciousness. Taking a parts view is useful in many ways when one is investigating how phenomenal unity might break down, as for example in Chapters Six to Nine where he evaluates putative empirical cases of exceptions to the PUT. This usefulness argues for the wisdom of thinking of component experiences as token phenomenal experiences in their own right, and therefore, a mereological approach is clearly preferable to a Jamesian one-experience view (p. 23). But Bayne’s conclusion is not the only one available from the data. The fact that there are auditory and gustatory experiences within a single stream of consciousness is just as compatible with a single experience containing a multiplicity of *contents* as it is with a mereological view with multiple *experiences*. In fact, a multiplicity of contents each with “distinct causal profiles and a certain degree of functional autonomy” is just as satisfying an explanation of articulation as Bayne’s mereology of token experiences, so this argument does little to support the superiority of a mereological view over a one-experience view.

Turning to the second question, Bayne’s mereology—however useful—brings with it a little baggage. Mereology (from the Greek μέρος, ‘part’) is a term coined by Leśniewski in 1927 that denotes “the theory of parthood relations: of the relations of part to whole and the relations of part to part within a whole” (Varzi, 2015). The study of mereology is a complex matter indeed, and even something as simple as the word “part” admits a bewildering variety of interpretations and uses. In choosing a mereological approach to phenomenal unity, Bayne opens his theory up

to some rather complex and unsettled metaphysical issues, only a very few of which I will attempt to address here.

One of these issues is the idea that parts in mereology—whatever kind of metaphysical entity they happen to comprise—are usually capable of independent existence. I do not wish here to get involved in the ongoing debates over metaphysical entities, for example, whether there is a real distinction between universals and particulars or whether universals are real entities.¹⁹ I merely wish to establish the point that on most metaphysical systems, such as that of Lowe (2002) for example, metaphysical entities seem to fall into just two groups. On the one hand there are entities that have no parts at all; *universals* like redness or betweenness. On the other hand there are entities that have completely detachable parts, parts capable of independent existence: *abstract particulars* (sets and propositions); *events* (e.g. explosions); *substances* (e.g. animals) and *non-substances* (e.g. shadows). The parts may not be the same kind of thing as the whole in some cases, but they do not depend on the whole for their existence, such that if the whole did not exist, the part could not exist.

In the taxonomy of life example of subsumption, the parts seem metaphysically independent of the whole: not merely *capable* of existing without the whole, but *in fact* giving existence to the whole. It is the concrete particulars of individual organisms that exist in the physical world, while species, orders and classes are merely ways of grouping (describing) those individual organisms. They are abstract concepts we use to think about the concrete particulars that are the

¹⁹ See for example, Rodriguez-Pereyra (2014) and MacBride (2005).

organisms themselves, and they would have little meaning if the concrete particulars did not exist. The concrete particulars are completely detachable from the subsuming entity. But this is not the case with Bayne's component experiences, which his holism insists cannot exist independently of the whole. In fact, on the face of it, his parts look as though they are unique among metaphysical entities in this respect, with no precedent either within common metaphysical systems, nor in paradigmatic examples of subsumption like the taxonomy of life. Either a new metaphysical category needs to be created to accommodate Bayne's parts, or there is something wrong with his account, or this naïve interpretation of his account is at fault. The first option seems rather drastic, although given the strangeness of phenomenal consciousness generally it is a door best kept open.²⁰ I consider the latter two options in the discussion that follows.

2.3. Bayne's "parts"

The third question posed above was about the exact nature of the "parts" in Bayne's account. In order to understand what exactly Bayne's parts might be, and to assess their validity, I turn to the distinction between *experiential parts* (EP) and *no experiential parts* (NEP) views of phenomenal unity (Brook & Raymont, 2014; Lee, 2014a; Raymont & Brook, 2009). After defining EP/NEP I will compare it to Bayne's account of atomism/holism. In this way I identify the problem at the heart of the tension between Bayne's mereology and his holism (the fourth question).

²⁰ For example, if Schwitzgebel (2014) is right, then phenomenal consciousness may well warrant a new metaphysical category.

In brief, EP views see a subject's total phenomenally conscious experience at a moment in time as being comprised of smaller, individual experiences, each of which is itself phenomenally conscious and could exist as a phenomenally conscious state independently of the total state of which it is part (i.e. detachable). Hume's bundle theory (Hume, 1739) can be seen as its ancestor. NEP views see a subject's total phenomenally conscious experience at a moment in time as being just one single experience that is *not* comprised of smaller, individual experiences, each of which is itself phenomenally conscious and capable of existence independently of the whole state (i.e. non-detachable). Brook and Raymont (2014) usefully clarify this in notation: if E is a phenomenal experience of intentional object ' o ', and a subject simultaneously has a unified experience of two intentional objects, o_1 and o_2 , then on the EP view, this experience may be represented as $E(o_1) \& E(o_2)$ whereas on the NEP view, it will be represented as $E(o_1, o_2)$. Further, Brook and Raymont point out that the proponent of EP can consider the view a characterisation or description rather than an explanation of phenomenal unity. But of course descriptions are part of the foundation for explanation and will inevitably influence or constrain the kinds of explanation available. Thus it is important to find the most accurate description available if one wishes to explain.

There is some confusion in the literature as to where Bayne's account stands in relation to the EP/NEP distinction. I give two examples here. Brook and Raymont's considered opinion is Bayne's subsumption relation is clearly an EP view. Yet even they find themselves equivocating at times. For example, referring to Bayne and Chalmers (2003) they state at one point,

“how their view stacks up with respect to EP is not entirely clear”

(Brook & Raymont, 2014, sec. 6.1).

But elsewhere they are unequivocal on the matter. They seem to assume that if subsumption requires that there be parts, those parts can only be EP parts.

“One feature of subsumption is that it requires that there be experiential parts. Thus those who wish a characterization of unified consciousness to be neutral on this issue will look for a different one” (Brook & Raymont, 2014, sec. 2.1).

“clearly this approach would require that there be experiential parts” (Raymont & Brook, 2009, p. 566).

They are not alone in finding Bayne’s account confusing on this point. Masrour also has trouble pinpointing exactly how Bayne wants us to think of the component experiences. He describes his own distinction between Newtonian and Leibnizian models of unity (Masrour, 2014a, 2014b). Briefly this distinction revolves around what grounds the unity of consciousness, a question we shall shortly see is crucial to understanding Bayne’s parts. *Newtonian* views hold that experiences are unified in the way that points in the same geometrical space are unified; that is in virtue of being in the same single space. This view is strikingly similar to Searle’s *unified field theory* which is paradigmatically NEP (Brook & Raymont, 2014). *Leibnizian* views (Masrour’s own preference) ground phenomenal unity in the connections between experiences such that the whole experience emerges from the

connections between the parts that comprise it, just as unified space emerges from the connections between the points that comprise it. Note the similarity to the taxonomy of life example, suggesting that subsumption is generally an EP kind of relation. I will not delve any further into Masrour's distinction other than to point out the interesting fact that on his own reading, as capable a mind as Masrour is unclear as to whether Bayne's account is Newtonian or Leibnizian. While Bayne confirms it is Newtonian in his response to Masrour (Bayne, 2014a) he goes on to say that there is still something to be said in favour of Leibnizian views. This might be Bayne being polite, but I believe it highlights a lingering ambiguity in Bayne's account that I tease out below.

So far as I know, Bayne himself has not addressed the EP/NEP distinction directly in print. He occasionally uses the term "experiential parts" (Bayne, 2005, p. 503, 2010, pp. 21, 36) but in a much more general sense than Brook and Raymont's technically defined usage. Nonetheless, we can find some clues as to what he might say on the matter, for he makes two significant and somewhat central points in *The Unity of Consciousness*. First, and possibly in apparent contradiction of Brook and Raymont's second quote above, he insists that his subsumptive mereological conception is neutral on the atomism/holism distinction. Second, he argues strongly on other grounds for holism, which is most naturally thought of as an NEP view (but see below). On the first point he writes,

"the mereological account ... is strictly neutral on the question of whether in seeking to understand the structure of consciousness we should assign explanatory priority to the parts (Searle's

‘building block’ model) or to the whole (his ‘unified field’ model)”
(p. 36).

What is interesting about this passage is that the examples Bayne uses are in fact paradigmatic of the EP (‘building block’ model) and NEP (‘unified field’ model) distinction. Perhaps we should then take it that Bayne holds his mereological conception to be neutral not only on the atomism/holism distinction but also on the EP/NEP distinction? I believe this is probably his position, although he nowhere says so explicitly. Second, in spite of this neutrality, he argues on *other* grounds strongly in favour of holism. But does this necessarily entail a similar commitment to NEP? The interpretations of Brook & Raymont and of Masrour above have not fully clarified the answer to this question. In order to better answer it, and to therefore get to the issues that underlie the apparent tension in Bayne’s account, I now turn to a more detailed comparison between NEP and Bayne’s holism.

In the table below I compare what I take to be the relevant features to my argument of Brook and Raymont’s EP/NEP distinction and Bayne’s version of the atomism/holism distinction. While the two distinctions mostly overlap, there are two significant questions on which they diverge.

Table 1 Two distinctions about phenomenal unity. Points of divergence between the two distinctions in bold.

	EP/NEP Brook & Raymont (2014)	Atomism/Holism Bayne (2010)
Which is metaphysically prior, component experiences or total experiences?	EP says component experiences are metaphysically prior. NEP says total experiences are metaphysically prior.	Atomism says component experiences are metaphysically prior. Holism says total experiences are metaphysically prior (pp. 36, 225-226).
Are there discreet and independently phenomenal component experiences?	EP says yes. NEP says no: there are no parts, only a single whole experience.	Atomism says yes. Holism says no: there are parts but they are not detachable (p. 225).
Is phenomenal unity grounded in a relation between component experiences?	EP says yes: that is all phenomenal unity is. NEP says no: phenomenal unity is grounded in there being only a single total phenomenal state, albeit with a multiplicity of something else, e.g. contents. ²²	Atomism says yes: that is all phenomenal unity is. Holism can also say yes: it can allow that phenomenal unity is grounded in a relationship (e.g. of subsumption) between component experiences and the whole. ²³
Is the relation between unified components an internal or an external relation?	EP says an external relation. NEP says an internal relation.	Atomism says an external relation. Holism says an internal relation (p. 236). ²⁴
Should we employ a bottom-up methodology or a top-down to explain phenomenal unity?	EP says bottom-up. NEP says top-down.	Atomism says bottom-up. Holism says top-down (p. 226).

²² Brook and Raymont cite Tye as a paradigmatic proponent of NEP. He says, “sense-specific experiences do not exist. They are a figment of philosophers’ and psychologists’ imaginations. And there is no problem, thus, of unifying these experiences. There are no experiences to be unified. Likewise within each sense” (Tye, 2003, p. 28).

²³ For example, “Experiences are phenomenally unified with each other exactly when they exemplify conjoint what-it’s-likeness. When a pain and a thought are phenomenally unified with each other the subject in question will enjoy the phenomenal character associated with the pain and that which is associated with the thought, and these two phenomenal characters will be experienced ‘together’” (Bayne, 2014a, p. 522).

²⁴ “Holists, of course, have no need to posit a mechanism that is specifically responsible for phenomenal binding, for the holist holds that components of a subject’s total phenomenal state are brought into being as the constituents of that state.”

Our fourth question above asked where the tension in Bayne's account might lie. I believe it lies in these two points on which his holism diverges from NEP. The first point of divergence is that where NEP holds there are no parts of the whole experience that are experiences themselves, Bayne's holism agrees with EP that there are parts, but unlike EP, insists that they are non-detachable, at least as token phenomenal experiences. Call this the *parts question*. On this question, Bayne's holism is neither pure EP nor pure NEP. The second difference is that while NEP grounds phenomenal unity in the singleness of experience, Bayne's holism agrees with EP in grounding phenomenal unity in a mereological relation of subsumption between the whole and the parts.²⁵ Call this the *grounding question*. These two questions suggest that Bayne's holism is neither EP nor NEP, reflecting the metaphysical strangeness of his parts that I discussed in Section 2.2. His account of holism in Chapter Ten reflects this incongruence between his holism and NEP, where his argument revolves around the analogy of "The Quilt of Consciousness" (also the title of the chapter).

"our approach to consciousness should be holistic ... One can identify particular fine-grained states of consciousness within this overarching state—just as one can pick out particular squares

²⁵ Does Bayne claim to provide a grounding for phenomenal unity? "Although my appeals to mereology were primarily designed to illuminate the notion of phenomenal unity rather than to point to some independent relation that might function as its metaphysical basis, I am inclined to think that a case can be made for grounding phenomenal unity in mereological relations" (Bayne, 2014a, p. 524).

within a quilt—but these states are not independently conscious ...
the structure of consciousness is fundamentally holistic” (pp. 244,
248).

The analogy of squares in a quilt is meant to relay holism, but surely it is a much more apt analogy for atomism and EP? How does a quilt of squares sewn together differ from a wall of bricks cemented together? And of course, the latter is Searle’s paradigmatically atomistic and EP *building block theory*, from which Bayne takes some trouble to distance his own view (pp. 35-36). If we are to make sense of Bayne’s view and find a resolution for the tension between mereology and holism, we need to understand what he means by discrete experiences that are nonetheless *not* independent units of consciousness.

2.4. Can the tension be resolved?

Bayne’s mereology defines the parts of experience as being themselves token phenomenal experiences on the basis of the tripartite approach to individuating experiences (subject, time, phenomenal character). But on his holism, these parts have some unique features: they are non-detachable from, and metaphysically subsequent to the whole experience, casting into doubt that they are genuine parts at all.

In this section I consider two ways one might try to resolve this tension. The first way involves seeing Bayne’s account as being more a way of speaking than an

actual account of the structure of phenomenally unified consciousness. Here, the term “parts” is merely instrumental—employed for a particular purpose without committing to any specific way of defining the parts as genuine ontological entities. But while this would evade the tension in his account, it would cost him much in terms of the usefulness of the account, and his views on the nature of the differences between his account and Tye’s one-experience view make this interpretation untenable. Second, I consider an interpretation that accepts the parts as genuine ontological entities and I try to pin down exactly what Bayne could mean when he speaks of “formally decomposing” a total phenomenal state into its parts. One understanding in particular promises to resolve the tension, the one that depends on attributing phenomenal unity to only some elements of a token component phenomenal state but not others. However, this leads not to Bayne’s holism, but to an NEP view.

2.4.1. *Verbal interpretation*

One way to resolve the tension between Bayne’s mereology and his holism might be to consider his talk of parts of experience as merely a useful verbal tool for investigating phenomenal unity, without saying anything about whether such parts are true ontological entities, somewhat like our talk of taxonomic classes and orders, which are just convenient ways to group living organisms. This approach accords with some of his motivations for a mereological approach discussed above (Section 2.2): it respects the distinct causal profiles of component experiences and makes it easier to handle empirical cases like split-brain patients. This is

something that Bayne himself has considered (recall Bayne casts his mereology as falling under “the received view”).

“The foregoing threatens to problematize the supposed contrast between Tye’s one-experience view and the received view, for Tye and the proponent of the standard view may simply be operating with different conceptions of experience. Perhaps the dispute between them is merely verbal. Perhaps the substantive disagreement—if there is one—concerns the question of how experiences ought to be individuated: should they be individuated in terms of instances of phenomenal properties (give or take a bit), or in functional terms? But is this really a substantive issue? There is, I think, a case to be made for thinking that there is no one privileged way of individuating experiences. Perhaps one ought to individuate experiences depending on one’s theoretical needs, and different theoretical needs might require different ways of individuating experiences. If this is right, then it might be argued that the one-experience model and the received view need not be viewed as competitors” (Bayne, 2005, p. 498).

“The notion of a token experience is elastic, and different approaches to the individuation of experiences might be appropriate in different contexts” (Bayne, 2010, p. 24).

Therefore, we might be justified in thinking of Bayne's parts as merely instrumental.

"One might think of subsumption as analogous to a sort of mereological part/whole relation among phenomenal states, although this should be taken as an aid to intuition rather than as a serious ontological proposal, at least at this point" (Bayne & Chalmers, 2003, p. 40).

This is also how Lee understands Bayne.

"parts of an experience are mere abstractions from the whole, rather than independently existing states" (Lee, 2014c, p. 766).

To summarise, this interpretation resolves the tension between Bayne's mereology and his holism by considering talk of "component experiences" to be merely instrumental. If Bayne sacrifices any kind of commitment to the ontological existence of independent component experiences then he can maintain the metaphysical and epistemic priority of the whole as required of any holistic view. His parts are non-detachable because they don't really exist as true concrete particulars. As a bonus, the apparent conflict between mereological views and one-experience views turns out to be merely verbal. It all comes down to how we define the term "experience". Call this the *verbal interpretation*.

Although there is some merit to this way of resolving the tension in Bayne's account, I believe it is not a satisfactory solution. It seems irreconcilable with Bayne's arguments in Chapter Two, particularly with his strong critique of Tye's one-experience view (see Section 2.2). Surely, if talk of parts were merely instrumental, and if the two views "need not be viewed as competitors", there would be no need to argue so strongly against Tye. It certainly seems as though Bayne wants more from his mereology than a mere instrument: he wants to insist that component experiences must be considered as genuine token phenomenal experiences in their own right, with their own unique "causal profile" and "functional autonomy". He rejects the NEP view that there are no real parts (the parts question, Section 2.3). This certainly sounds like much more than a merely convenient way of grouping things, as would be the case were the parts merely instrumental.

2.4.2. *Decomposition interpretation*

The chief alternative to the linguistic interpretation involves taking seriously Bayne's view of the parts as being token phenomenal experiences that are the same kind of thing as the total phenomenal experience. If the parts are real entities in this sense, then what role do they play in explaining phenomenal unity (the grounding question)? Bayne cannot espouse any view where the whole can be explained as a combination of the parts, for that would be atomistic and therefore inconsistent with his holism. His mereology cannot be *that* kind of mereology and his subsumption cannot be exactly the kind of subsumption involved in the

taxonomy of life. He does argue though that the whole can nonetheless be decomposed, albeit in a special way.

“Holists can allow that the phenomenal field can be formally decomposed into discrete experiences, but they will deny that these elements are independent atoms or units of consciousness”
(Bayne, 2010, p. 225).

Into what kind of parts can a phenomenal field be formally decomposed? I argued in Section 2.2 that the very concept of non-detachable parts is problematic. Given that, and given that talk of parts makes Bayne’s account fall somewhat between the apparently mutually exclusive and exhaustive poles of EP and NEP, the wisdom of trying to marry mereology with holism looks less attractive. If the holist explanation of phenomenal unity is not grounded in a relationship between experiences, why have parts at all? Whatever explanatory need they are meant to fill can certainly be filled with less troublesome entities than non-detachable token phenomenal experiences. Brook’s words are relevant here.

“Paul Raymont and I [2009] have presented arguments against subsumption and in favour of the Jamesian one-state picture just sketched. Bayne discusses them [34–5] but may not have fully appreciated their force ... Even on the subsumption account, there has to be one experience whose unity does not consist in its being a part of a ‘bigger’ experience. So why posit a bunch of ‘smaller’ experiences, too? To be sure, we have to posit a multiplicity of

something—but a multiplicity of objects of experience and ways of being conscious would do everything needed” (Brook, 2012, p. 600).

Can we save Bayne’s non-detachable parts by finding a way to make sense of them and a valuable role for them to play in grounding phenomenal unity? Below I consider two ways we might try to do that using a decomposition interpretation of Bayne’s parts. The *flying man interpretation* begins with the hypothetical case of a total maximal state that has just one component and from there derives non-detachable parts, while the *lone wolf interpretation* begins by considering where phenomenal unity obtains within the structure of phenomenal experience. Both interpretations however fall short of making sense of non-detachable parts.

Flying Man Decomposition Interpretation

Bayne’s holism depends on a top-down approach to phenomenal unity—we begin with the whole, and from there explain the parts. One interesting way to do that is to adapt a thought experiment first proposed by Avicenna—the “flying man” (Marmura, 1986). He imagined a man created out of nothing, suspended in still air, and therefore with severely limited sensory input. A modern version of this might be the sensory deprivation chamber, where the subject is immersed in buoyant fluid with her ears, eyes, nose and mouth all occluded. Even here, interoceptive inputs would still remain, not to mention non-sensory cognitive states,²⁶ so for the purposes of this argument, imagine it is possible to shut down every single source

²⁶ Whether pure non-sensory thoughts, like “three is a number”, have a phenomenal character, is an open question. A good overview is Bayne and Montague (eds) (2011).

of sensory input, internal or external, and every other kind of phenomenal experience including thoughts, beliefs etc.—all except for just one simple phenomenal experience²⁷ (call it A_i for “isolated”). Perhaps this might be achieved by chemically ablating the relevant afferent nerves and regions of the brain. Surely there is no *prima facie* reason to doubt that A_i is a token phenomenal experience? And if it is so in this extreme hypothetical case, then why should it not be so in the normal case where it is part of a more complex total phenomenal experience? Yet it is also clear that even when it is alone, it does not exist independently of the total phenomenal state. In this way we have established that A_i is indeed a token phenomenal experience in its own right, while preserving its non-detachability. Imagine this process could have involved other experiences, B_i and C_i and so on. When all these experiences happen to occur simultaneously within a single total phenomenal experience, they are validly token experiences that are non-detachable parts of the whole.

There are at least two problems with this interpretation. First, we can draw an important distinction between A_i and A_c —the latter being the experience of A in the context of other experiences with which it is unified. Below (Section 3.1), we shall see that quite often the phenomenal character of a component experience changes due to influence from other phenomenal characters. This change means that the phenomenal characters of A_i and A_c are different, different enough to individuate them on Bayne’s tripartite conception as numerically non-identical phenomenal states. So even if A_i is a part of the total phenomenal state, that proves nothing about A_c , which is the realistic (as opposed to hypothetical) part in which

²⁷ For the purposes of my argument, just how small a “single” experience might be is not important.

we are really interested. Second, it could be argued that A_i is not really a part of the total state, but merely the total state made smaller. Analogies fail me here since all the analogies I can think of involve spatially extended entities that have detachable parts that can be detached to make the entity smaller, but that is not the case here. Perhaps the best (imperfect) analogy available is that of the expansion of the universe from the Big Bang until now. Space is not being ‘added’ to the universe—the existing space is being stretched bigger. The universe moments after the Big Bang was did not have less space in it—the same space was squeezed tight. Similarly here A_i might be thought of not as being a total experience with fewer parts so much as a total experience squeezed tight. And if that is so, then we have not really established that A_i is a part in Bayne’s mereological sense at all. Taking it in the other direction, the addition of other ‘parts’ does not involve building block-style addition, but ‘stretching’ the existing experience to cover more contents. Tye could argue that we never had anything but a single experience here, just with more or less complexity in its contents. So, the flying man interpretation does little to make sense of non-detachable parts, for the ‘part’ here is not really a part at all, but just a shrunken whole. What is more, turns out to fit equally well with NEP views rather than favouring Bayne’s holism.

Lone Wolf Decomposition Interpretation

The second possible deployment of the decomposition interpretation is to take Bayne’s talk of non-detachable parts as meaning that a phenomenal state A that is detached from the total phenomenal state T loses its *phenomenality*, and therefore ceases to exist *as* a token *phenomenal* state. A lone wolf does not have the property

of being capable of killing a caribou, but it acquires that property only as a member of a coordinated pack.²⁸ This is a little like the concept of emergence: the parts acquire novel properties—in this case phenomenality—only when they combine as a whole in certain ways. Detach the part from the whole and it loses those emergent properties, and therefore becomes a different thing. This not only gives us non-detachable parts but, contra NEP, also grounds phenomenal unity in the relation between the parts and the whole, since *A* is phenomenally unified with other component experiences *only* when it bears the relation of being subsumed by *T* just as Bayne’s holism requires.

Unfortunately, this interpretation, when taken to its logical conclusion, also fails to resolve the tension between mereology and holism, since it collapses back into an NEP view, as I will now argue. Let us maintain the notation A_c for a component experience that is phenomenally unified with other component experiences within *T*, and add A_n for exactly the same mental state but without the property of being phenomenally conscious. Thus for example, if mental state *A* has as its contents the famous four-note dramatic stanza from Beethoven’s fifth symphony, “GGGE *b*”, then there is a mental state A_n with the content in musical notes, “GGGE *b*” that is not phenomenal—say, it reaches your ears while you sleep and it makes you jolt, but doesn’t actually wake you up. If on the other hand you were to hear “GGGE *b*” while awake and were phenomenally conscious of it, the same content would now be the phenomenal mental state, A_c . Many things about A_c and A_n are identical: the sensory content, perhaps even most of the neural correlates of the mental state,

²⁸ Thanks to Rachael Brown for this idea.

and so on. The only difference between the two mental states is whatever it is that is the difference between a phenomenal and a non-phenomenal mental state: phenomenality, whatever that might be.

It is here at the level of phenomenality that we must look for the key to phenomenal unity. So what is the difference between A_c and A_n ? What is it that makes one phenomenal and not the other? This is the deeper mystery of the nature of phenomenal consciousness itself, which is of course far too great a topic for this thesis to try to solve. But fortunately we don't have to solve it here—we just need to know what Bayne thinks is the difference between the two states. For Bayne, “*GGGE b*” acquires phenomenality just when it “enters” a phenomenal field, and loses phenomenality, and therefore also its identity as a token *phenomenal* state, just when it leaves that field.

“Holists, by contrast, hold that the components of the phenomenal field are conscious only as components of that field ... [Atomists recommend that] we should focus on understanding the mechanisms responsible for generating the atoms of consciousness. Holists, by contrast ... will doubt whether we can understand consciousness by focusing on the components of the phenomenal field in a piecemeal manner. Instead ... we should look for the mechanisms implicated in the construction of the entire phenomenal field” (pp. 225-226).

Bayne's holism dictates that phenomenal unity has a top-down mechanism. Phenomenality itself is something that is created as a whole, not in pieces that are then bound together. The reason A_c loses its phenomenality when it becomes A_n is not that it has changed intrinsically, nor even that it has changed its relations to other component states, but that it has ceased to participate in that top-down process that bestows phenomenality on all the mental states that are currently components of T . This view makes better sense of the empirical fact that mental states are constantly "joining" and "leaving" T as time passes: it is not that they each have their own individual mechanism of phenomenality that is being switched on or off, but that they participate or withdraw from participation in that single mechanism that produces phenomenality and bestows it on certain mental states *en masse* (p. 244).

This aspect of Bayne's account turns out to exacerbate the tension in his account by supporting the NEP answers to the grounding and parts questions rather than the answers given by Bayne's holism. On the parts question: the field that bestows the property of phenomenality is single, and mental states that enter that field are not proper parts *of* the field, but are another kind of entity that is affected *by* the field. This favours the NEP answer to the parts question: there are no parts as such, only a single whole experience. There is a mereology of many things—contents and neural substrates and suchlike—all of which are non-phenomenal. But there is no mereology of phenomenality itself, which is a single mechanism or field acting on multiple mental states. And therefore, on the grounding question: the *unity* of the experience is grounded in the singleness of the field and not in the relation between the field and the component mental states. The *phenomenality* of each

component experience is grounded in their relationship of participation in that field, but the *unity* of experience is not grounded in this relationship, but in the singleness of the field. This favours the NEP answer to the grounding question: phenomenal unity is grounded in there being only a single total phenomenal state, albeit with a multiplicity of something else (in this case, token mental states or contents). So the lone wolf decomposition interpretation finds Bayne's account collapsing back into an NEP account. I will develop this topic further below in Section 3.3.

2.5. Summary

In this chapter I began by refining our understanding of Bayne's mereology and of his particular brand of holism. I then introduced the EP/NEP distinction and used it to pinpoint the apparent tension in Bayne's account between his mereology and his holism. At the heart of this tension lies the exact nature of his parts of experience. To try to understand his parts better, I compared Bayne's holism to NEP views and found them to diverge on two chief points: the parts question and the grounding question. I then considered three interpretations that might resolve the apparent contradiction in the idea of parts that are non-detachable. The verbal interpretation failed because Bayne clearly means more when he talks of parts than just a convenient label, especially as evidenced by his emphatic disagreement with Tye's one-experience view. The flying man decomposition interpretation failed because its "parts" turned out not to be parts at all, thus favouring an NEP view over Bayne's holism. The lone wolf decomposition interpretation failed

because the structure of phenomenal unity it is built on leads again to an NEP answer to the parts and grounding questions. So far the tension between Bayne's mereology and his holism appears unresolved, and the coherence of his attempt to synthesise the two together in a single model remains in doubt.

There is another piece to the puzzle that initially holds some promise of resolving the tension, but turns out instead to increase this tension rather than ease it—the idea that parts intimately influence each other's phenomenal character. That is the topic of the next chapter in which I explore this feature, and use it to tie together some loose ends.

3. Phenomenal interdependence

Adding to the pressure on Bayne's account is the idea that component experiences are not completely discrete entities, but sometimes seem to melt into each other in such a way that they influence each other's phenomenal character. In this chapter I elaborate a little on this concept by introducing Dainton's concept of *phenomenal interdependence* (PI) and teasing out some of its nuances. I then draw some specific implications for Bayne's account, particularly on the grounding and parts questions (Section 2.3). I argue that the indistinct borders between component experiences are more compatible with the NEP answers to the grounding question than with the answer given by Bayne's holism, in that they ground phenomenal unity in the singleness of the total phenomenal state, rather than in mereological relations. On the parts question, PI is unable to justify his problematic concept of non-detachable parts (Section 2.2). I conclude that this mutual influence between component experiences—something Bayne endorses with conviction—actually

adds to the tension between Bayne's mereology and his holism rather than resolving it.

3.1. Phenomenal Interdependence and its 'strengths'

The idea that the parts of experience influence each other's phenomenal character may be traced back at least to William James and an obscure Rev. James Wills (Dainton, 2000). Here, I will employ the term *phenomenal interdependence* (PI) (Dainton, 2000, p. 191) to denote this idea. Think of a token component phenomenal experience A . As above, we can distinguish two versions of A depending on whether it is experienced alone in isolation from any other experiences or as part of a complex unified phenomenal experience, A_i (for in *isolation*) and A_c (for in *context*). PI is the idea that when states A_c and B_c are experienced together, the phenomenal character of A_c influences that of B_c and vice versa. It is this difference in phenomenal character that distinguishes A_i from A_c . Since PI involves *influence* between phenomenal states that results in the *integration* of those states, I will occasionally use those terms to refer to PI.

An example of PI in action is de Vignemont's distinction between *additive binding* and *integrative binding*.

“[Additive binding of phenomenal states] results from the combination of sensory information that is not redundant because it is about distinct features of distinct parts of the object. However,

there are many cases in which one modality influences another modality in one way or another. In particular, integrative binding results from the fusion of sensory data that is redundant” (de Vignemont, 2014b, p. 130).


de Vignemont is thinking here of multimodal fusion, say, hearing a word spoken and seeing lips forming that word. Because there is a redundancy of information—the word itself being represented twice—the two modal experiences fuse together into a single experience. This is one way that phenomenal states can influence each other’s character, but there are many others.²⁹ There is good evidence that PI applies not just between modalities, but within modalities as well. For example, within vision, Anderson and Winawer (2005) discuss some captivating illusions of changing perception of lightness³⁰ (Figure 2, below). In these cases, if A_i is the experience of the moons alone without any sort of background (*intrinsic phenomenal character*) and A_c is the experience of the moon against a textured background that makes it seem either lighter or darker than A_i (*contextual phenomenal character*), then by Bayne’s tripartite approach to individuation of experiences, A_i and A_c are numerically different token phenomenal experiences of

²⁹ For example, see Bennett & Hill (2014b), De Gelder & Bertelson (2003) and Foxe & Schroeder (2005).

³⁰ There are many more fascinating illusions showcased in the *Illusion of the Year Awards*, <http://illusionoftheyear.com> (accessed April 2015). Unlike the famous Müller-Lyer illusions these Anderson and Winawer illusions do seem penetrable. That is, once you have understood what is happening, it seems possible to voluntarily overcome the illusion and see the moons as identical, in spite of the differing backgrounds. This difference between the two illusions is evidence that they are caused by processes going on below the level of phenomenal consciousness (see below), but in one case those processes can be controlled consciously, while in the other they cannot.

the same moons, and the differences between them arise from the different contexts in which they are experienced.

a



b



Figure 2 Visual Illusion due to Anderson and Winawer (2005, p. 80)—see text for discussion. The two sets of both moons and chess pieces are intrinsically identical, but have different phenomenal characters depending on the background.

Is PI a universal feature of phenomenal consciousness, or is does it only occur sometimes? Below, I am going to argue that PI exacerbates the tension between mereology and his holism in Bayne's account, so the extent to which PI obtains is quite pertinent to my argument. Ultimately, Bayne's project is to provide an account that promotes our understanding of the nature of phenomenal unity. If PI is ubiquitous, and component experiences always influence each other, then it is quite possibly an intrinsic and integral feature of phenomenal unity, and any account of phenomenal unity that does not incorporate it runs the risk of being seriously deficient. If on the other hand, PI is the exception instead of the rule, and many component experiences are completely uninfluenced by other component experiences with which they are phenomenally unified, it is more likely to be a peripheral feature of phenomenal unity, and it may not be so important that Bayne's account does not incorporate it sufficiently. Dainton's (2000, pp. 191–192) "strengths" of PI,³¹ which I summarise below, provide a useful basis for this discussion.

PPI – *Partial Phenomenal Interdependence*

PI applies to only some parts of total experiences, or to all parts of only some total experiences.

³¹ Koksvik (2014) covers similar ground to Dainton but takes a different approach. He proposes three models for the degree of interdependence between component experiences, but his chief interest is more about the contribution of parts to the character of the whole, than specifically about the effects of the interdependence on the components themselves. There is of course significant overlap between his discussion and Dainton's.

CPI – *Complete Phenomenal Interdependence*

PI applies to all parts of all total experiences.

Strong CPI

E.g. the very implausible *holographic conception* of experience, where the entire and exact character of a total experience is encoded in each of its parts.

Weak CPI

Each part of a total experience would be different in some way if any of the other parts were different in character.

To Dainton's list we can add one more item for completeness.

NPI – *No Phenomenal Interdependence*

PI never occurs.

There are two other brief but significant points about how I employ PI here that are worth clarifying. First, PI is not *just* Gestalt unity, where "the whole is greater than the sum of its parts". Although that may also happen to obtain, I am more focused here on the changes to the phenomenal character of component experiences themselves, rather than on the phenomenal character of the total experience as a whole. Second, I am interested in a particular sort of mutual influence between component experiences, the kind of influence that occurs within the closed system of a single organism. This must be distinguished from other kinds of mutual influence, say the kind that occurs when two people engaged in an

intricate argument are sharing ideas, influencing each other's ideas, experiences and the phenomenal character of each other's mental states. Here, there is no co-consciousness between their phenomenal states. At most, there may be some 'empathy' between them, in the sense that they may mimic or imitate each other's phenomenal states by running 'simulation software' in their heads, but this does not make their token phenomenal states in any way phenomenally unified, and the underlying mechanism of this influence is likely to be significantly different.³² Armed with this nuanced concept of PI we can begin to apply it to Bayne's account.

How prevalent is PI, where does Bayne stand on this question, and how does it affect my argument? I address each question in turn. First, examples like the Anderson-Winawer illusions strongly suggest that NPI is false and at least PPI is true, and few would disagree with this.³³ On the other hand, Dainton argues that extreme forms of strong CPI, such as the *holographic conception* of experience,³⁴ are implausible, and I tend to agree. There is little evidence from introspection to support it and the complexity of the neural connections likely to be necessary to underwrite it seems prohibitive. However, weak CPI in certain forms may be neurally plausible if the interdependence occurs only very late in the processing streams. Dainton however, considers any form of CPI highly implausible.

³² Recent interest in extended cognition may have more to say about this fascinating topic than is, alas, beyond the scope of this study. A good general overview is Menary (2010). For more on mirroring and empathy see for example, Schütz-Bosbach et al (2006). Work has also been done on the neural basis of this interaction between minds, for example, (Frith & Frith, 1999).

³³ Koksvik (2014, p. 112) cites Frank Jackson as denying PI altogether (supporting a *no-context view*) "in conversation" although not yet in print.

³⁴ "[T]he entire and exact character of a total experience is in some manner reflected or encoded in each of its parts" (Dainton, 2000, pp. 191–192).

“But on the face of it, it seems unlikely that small differences within a single perceptual field would always or even usually impinge on the local characteristics of the remainder of the experience within the same field. If the pile of books to my left were differently located, a centimetre or so further to the right, say, I see no reason to suppose the contents of the remainder of my visual field would in any way be different” (Dainton, 2000, p. 195).

This argument however does not discount CPI altogether as Dainton suggests. If the interdependence of parts is allowed to vary in intensity—anywhere from nearly negligible to utterly transformative—then we ought not be surprised if the nearly negligible interdependence escapes our casual introspective scrutiny, and appears often to be absent therefrom. It may require special experiments to reveal its presence (de Vignemont, 2014a), or given the ways our brains work, it may be impossible to reveal it at all, even if it occurs and has real effects on our behaviour. And there may be many subtle ways in which the interdependence may occur quite commonly, that are revealed to introspection only when we specifically look for them. For example, think of Dainton’s hypothetical pile of books above. The phenomenal character of a cup of tea located near them may be ever so slightly different if they were one centimetre closer to it. At their original distance, their toppling over may present no risk of hitting the cup and spilling it, but that one centimetre closer may be enough to now make the books a real risk of doing so, with disastrous results. This difference may be enough to make the cup and its

spatial relations much more prominent in my attentional field, thereby changing its phenomenal character. As Merleau-Ponty (1962) pointed out, our perceptions are more than just integrated sense *data*. The data is further integrated or ‘charged’ with layers of meaning.

Second, on the question of where Bayne stands on PI, we find that he argues quite strongly for PI *between* modalities. He writes of the extensive and quite impressive evidence for interdependence between the processes of different modalities, even quite early in their processing streams. Bayne’s arguments for inter-modal PI are not arguments for CPI, as is evident from the quotes below, where he uses words like “essentially” rather than “necessarily”, and “often” rather than “always”, and speaks of rules having exceptions. But it is clear that he is convinced of the truth of at least PPI, and I believe the empirical evidence he presents is persuasive.

“the ubiquity of multisensory interaction in the brain has led some researchers to describe the cortex as “essentially multisensory”

(Foxy & Schroeder, 2005; Ghazanfar & Schroeder, 2006) ...

Research extending back for over a century has revealed that perceptual processing in one modality is often modulated by perceptual processing in other modalities (Bertelson, 1988; Calvert, Stein, & Spence, 2004; De Gelder & Bertelson, 2003)”
(Bayne, 2014b, p. 21).

“Multisensory integration ‘is the rule and the not the exception in perception’ (Shimojo & Shams, 2001, p. 505)” (Bayne, 2014b, p. 30).

What of interdependence within modalities? The moon illusion above demonstrates that PI obtains at least sometimes within the modality of visual processing, and it would not be hard to establish it for other modalities (e.g. think of how gourmets mix flavours of foods and drinks). Bayne argues that the processes that integrate modalities together are likely the same kinds of processes as those that integrate within modalities. Therefore, contra Dainton, if inter-modal PI is widespread, so also will intra-modal PI be. The quote below refers to binding rather than interdependence per se, and I believe Bayne uses the word “integration” in a loose sense, rather than in the technical sense I am thinking of here. While binding and interdependence are certainly not co-extensive, I believe the principle involved may be reasonably taken to operate over both binding and interdependence. If interdependence often obtains inter-modally, then given the similarities in the mechanisms of inter-modal and intra-modal binding, it is reasonable to assume that interdependence often obtains intra-modally.

“Insofar as we understand the neural basis of consciousness, it would seem that a deep contrast between inter-modal integration and intra-modal integration is rather difficult to draw. There is no reason to suppose that the mechanisms of binding the output of *different* sensory modalities into a unitary experience of multi-modal objects, are, in essence, any different from the mechanisms

that bind the output of various discriminations *within* a modality.

Just as the brain binds information derived from different modalities into (usually) coherent multi-modal representations of objects, so too it binds colour, spatial, figure-ground, texture and other discriminations into a (usually) coherent *visual* perception of an object.” (Bayne, 2000, pp. 251–252).³⁵

Third, on the effect the strength of PI has on my argument, is this enough to ground my assertion that any account of phenomenal unity needs to incorporate PI? I believe that a good case can be made for weak CPI, but that is not essential for the purposes of my argument. Even if we grant that PI is not ubiquitous but just widespread, both intra-modally and inter-modally as Bayne seems to argue, the relevance of PI to the tension in Bayne’s account is maintained. If PI is part and parcel of much of our phenomenally unified experience, then anything it tells us about the nature of that phenomenal unity is going to count as important evidence in the evaluation of any account that claims to promote our understanding of phenomenal unity. It may even be essential to our understanding in the same way that a single case of time dilation overturns Newtonian physics in favour of Einsteinian relativity. Strengthening the value of PI in evaluating Bayne’s account is the fact that he himself has employed PI to evaluate a competing account of phenomenal unity, namely *atomism*, as follows.

³⁵ Italics in the original.

3.2. Bayne and phenomenal interdependence

PI has been understood by some to play little or no role in Bayne's account of phenomenal unity, to its disadvantage.

“On the Bayne-Chalmers view, multisensory integration is, on the face of it, less directly relevant to our understanding of the more general category of the unity of consciousness (where that includes subsumptive phenomenal unity) than it is on the Hill-Bennett view. Still, if one wants to understand total conscious states in terms of units of consciousness which are constitutively independent of the wholes which they comprise, then it may be that the best understanding of the basic units of consciousness is a multisensory understanding. In that case, multisensory integration will be relevant to a full understanding of total conscious states, in terms of what gets subsumed, even if not to the relation of subsumption” (Connolly, French, Gray, & Prettyman, 2011, p. 3).

But Bayne has deployed PI against atomism and against what he describes as “one manifestation of this atomistic approach to perceptual experience: the decomposition thesis” (Bayne, 2014b, p. 15), which holds that a total phenomenal state can be carved at the joints or decomposed into pure single-modality states. This sounds very much like a modality-based form of EP.

“The ubiquity of inter-modal integration puts further pressure on atomistic approaches to consciousness. Not only does it provide additional evidence against the view that perceptual features qualify as atoms of consciousness, it also tells against more moderate forms of atomism that conceive of the units of consciousness in modality-specific terms” (p. 235).

“[Inter-modal integration] indicates that the senses do not operate as a ‘mass of tributaries running in parallel’, as some atomists claim (O’Brien & Opie, 1998), but as a network of highly entangled and inter-dependent channels” (p. 247-248).

“the decomposition thesis struggles to accommodate the phenomenon of multi-sensory integration” (Bayne, 2014b, p. 30).

But Bayne’s deployment of PI against atomism has consequences for his own mereology. He argues that multi-sensory integration entails that the processes underlying phenomenal consciousness are more likely one big messy, entwined, multi-modal process than encapsulated, single-modality streams. I agree, and argue that this principle can and should be extended to include processing even within a single modality. If so, then just as Bayne argues that we should not think of phenomenal unity as carvable along modal boundaries, neither, I believe, should he think that it is carvable *within* modalities, along the borders marked by differences in phenomenal character. And if that is so, then this puts serious pressure on his mereology.

In support of this hypothesis, I put forward two arguments below. First, on the grounding question (from Section 2.3), I suggest that we must distinguish between interdependence and phenomenal unity, and that the most plausible mechanism underlying phenomenal consciousness suggests that they are different processes occurring at different stages. On extending the decomposition interpretation argument in Section 2.4.2, we find that the mechanism demanded by PI turns out to be one that grounds phenomenal unity in phenomenality itself as opposed to binding or interdependence, and therefore in a single experience, as per NEP, rather than in a mereological relation of subsumption between parts and wholes, as per Bayne's holism. Second, on the parts question (also from Section 2.3), I try to resolve the tension between Bayne's mereology and his holism using PI, but find little success.

3.3. On the grounding question

It is clear that PI is not identical to phenomenal unity. It is certainly conceivable (and Dainton argues, is most often the case) that mental states can be phenomenally unified without being in any way interdependent.³⁶ Also, even *non*-phenomenal mental states that are not phenomenally unified (because not phenomenal at all) still exhibit interdependence. Just think of the smoothness of

³⁶ Van Gulick (2014a) presents an argument for grounding phenomenal unity in representational integration and unity. This is not the place to critique it, but it has features that both support and count against the view I put in this paragraph.

the movement you make when you reach out your hand to pick up that cup full of hot tea by the pile of books. The fine adjustments to the commands being sent to the muscles in your arm are due to the content of the fine sensory data—visual, tactile and proprioceptive in this case—influencing the content of your motor commands, although perhaps the majority of both the afferent and efferent information, and the processing in between, never enters your field of phenomenal consciousness. So interdependence occurs between non-phenomenal mental states. What of phenomenally conscious mental states? Does the integration between them occur pre-phenomenally, or do states that are already phenomenal then influence each other? Call these two options *C-first* (consciousness-first) and *B-first* (binding-first) models of PI, represented in very rough thumbnail sketches in figure 3.

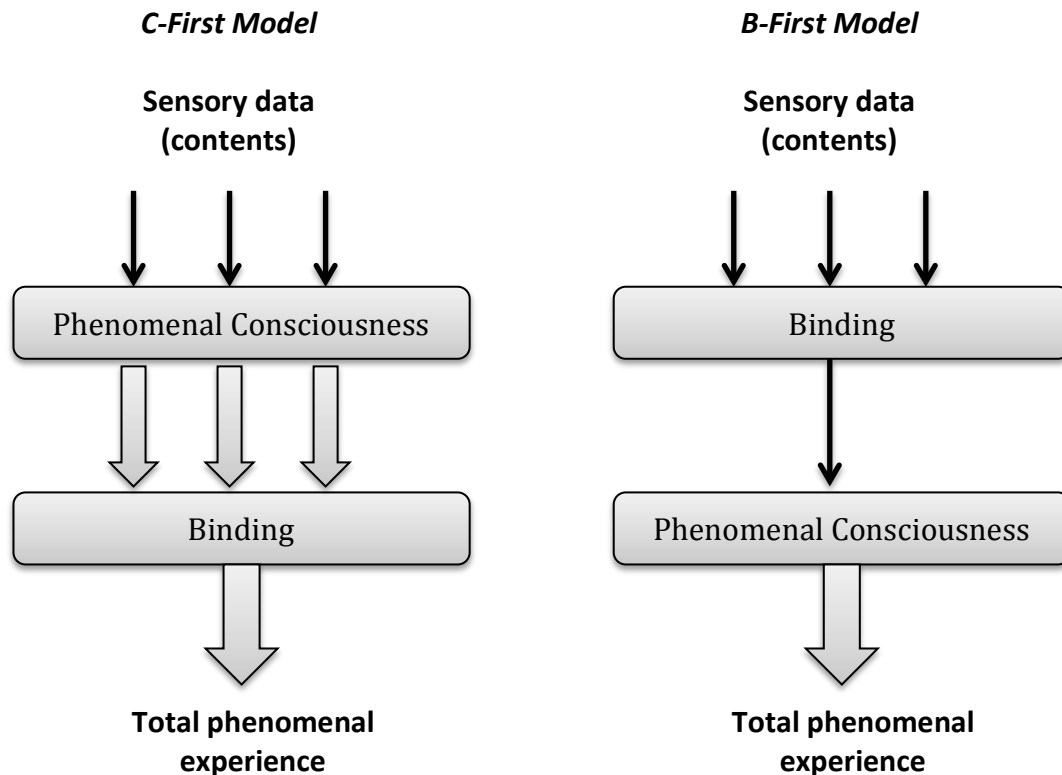


Figure 3 Two models of how sensory data becomes phenomenal experience. Line arrows signify non-phenomenal content while block arrows signify phenomenal content. Left: the C-first model where incompletely processed data acquires phenomenal character earlier in the process. Right: the B-first model where data acquires phenomenal character only after all the other processing, including binding, has already occurred.

On the *C-first model*, phenomenal character is endowed upon contents at some stage of processing *before* they are all bound into a single total phenomenal state. This entails that at least some component experiences possessing intrinsic phenomenal character will subsequently need to be bound together to create the total phenomenal experience. The C-first model is consistent with EP and atomistic views, but not with NEP or holistic views (including both Bayne and Tye). On the *B-first model*, all the binding occurs at pre-phenomenal stages of processing, and it is only when that is complete that a single non-phenomenal set of content is

imbued with phenomenal character as a single whole. The B-first model is consistent with NEP and holistic views (including both Bayne—I shall justify this shortly—and Tye), but not EP and atomistic views.³⁷

This suggests that it is important to distinguish *where* exactly within these models phenomenal unity obtains. But Bayne’s basic formulation of his view doesn’t do this. Instead, it seems to imply that phenomenal unity simply obtains over token phenomenal mental states *as such*.

“The mereological view treats phenomenal unity as a relation between token experiences—that is between particular mental states or events” (p. 21).

“when it comes to conscious thought, the frame of reference that concerns us must be that of the states themselves rather than their contents” (p. 18).

But when Bayne comes to discuss the concept of *articulation* (Bayne, 2005) we find there is more nuance in his view. Articulation, for Bayne, means that component experiences can combine in a variety of configurations (with some constraints). This ability to combine in different ways argues for the reality of

³⁷ Whether B-first or C-first correctly characterises human consciousness is a matter of empirical investigation that has not yet been satisfactorily resolved, so far as I know. However, there may be at least one strong theoretical argument against C-first. If the parts become phenomenal *before* they are bound together, why is it that we never *experience* them in this unbound state? What could it possibly mean that an unbound part is phenomenal yet not experienced by the subject? But a detailed critique of the two models lies well beyond the scope of this study.

parts rather than for a one-experience view (he seems to assume that this single experience must be homogenous).

“The intuitive explanation of why one can experience a dog barking without experiencing an itch and vice versa is that the experience of the dog barking is distinct from the experience of the itch. In other words, it seems natural to explain articulation in the *contents* of consciousness by appeal to articulation in the *vehicles* of those *contents*. I do not claim that phenomenal articulation provides a knockdown objection to the one-experience view, but it does, I think, provide a prima facie objection to it” (Bayne, 2005, p. 499).³⁸

Note how in developing his argument, Bayne admits that articulation as such occurs at the level of the *contents* or of the *vehicles* of consciousness, which like interdependence, can be pre-phenomenal (although they may acquire phenomenality). Articulated mental states become phenomenal *en masse*.

“The elements of a subject’s total phenomenal state do not enter consciousness as independent units but only *en masse* ... the components of a total state of consciousness come into being *as* the components of that total state. One need no more bind the various components of consciousness together than one need bind

³⁸ My italics.

the various organs of a body together—they come into being as unified with each other” (p. 244).³⁹

“there is good reason to think that perceptual features of most kinds are conscious only in the context of bound percepts. And if this is right, then radical forms of atomism ... must be false ... we should not think of feature binding as a process of ‘bringing different conscious experiences together’ (Bartels & Zeki, 1998, p. 2330), for there is in general no perceptual experience prior to feature binding” (pp. 232-233).

Putting together Bayne’s stated views on interdependence and binding, we can say that he sees non-phenomenal contents as being first bound together and integrated, and then as a whole, *en masse*, acquiring phenomenality. This concurs, of course, very well with his holism. But I want to emphasise here that on this model, the interdependence of the content of one mental state upon the content of another must clearly be occurring at the binding stage, which means it is pre-phenomenal. It is only the fully interdependent, fully entwined, fully bound contents that then become phenomenal “*en masse*”. None of the binding, thus none

³⁹ Italics in the original. I note again the difficulty in finding suitable analogies for Bayne’s holism. If Bayne in this quotation is referring to a process occurring through time, and if he means that the organs of the body are not created independently and then assembled together into a whole body (like Frankenstein’s monster) then of course he is right. But if he is seeking a reason for thinking the body is one, then of course something must bind the organs together in the sense of holding them physically together. This role is usually played by connective tissues, blood vessels, nerves and so on, as any anatomy student doing a dissection will confirm. In that sense, the organs of the body are detachable parts, unlike Bayne’s non-detachable parts of experience.

of the mutual influence, occurs *after* the contents acquire the property of phenomenality.⁴⁰ So Bayne's view of phenomenal unity is clearly a B-first model.

We come to an important difference between phenomenal interdependence and phenomenal unity. When non-phenomenal contents that are interdependent acquire the property of phenomenality *en masse*, their interdependence is of course preserved and mirrored in their phenomenal character, thus producing *phenomenal* interdependence. Phenomenal unity for Bayne, on the other hand, cannot *just* be the acquisition of the property of phenomenality by a mental state exhibiting some other kind of non-phenomenal unity. For one thing, he holds that the unity of consciousness cannot be grounded in any kind of unity (e.g. representational unity) other than phenomenal unity per se (Chapter Two). For another thing, when representationally unified content acquires the property of phenomenality, it does not thereby automatically become *phenomenally* unified. There is a kind of double dissociation here: representationally unified content can be phenomenally disunified (e.g. two scientists working on a theory together) and phenomenally unified content can be representationally disunified (e.g. split brain patients, see Chapter Nine). A similar argument can be mounted for other kinds of non-phenomenal unity.⁴¹ So unlike PI, phenomenal unity cannot be grounded in the phase of the B-first model that is pre-phenomenal, but must be grounded in the phase that is phenomenal. If phenomenal unity is to be understood as a relation at

⁴⁰ Incidentally, this means that Bayne's *prima facie* case from articulation against Tye's one-experience view quoted above does not hold up on closer examination. Tye can deflect it by positing a multiplicity of contents that can be articulated in different ways and produced by a multiplicity of vehicles pre-phenomenally, while grounding phenomenal unity as such in the singleness of the event in which those contents become phenomenal.

⁴¹ Such as object unity, subject unity and so on (Bayne & Chalmers, 2003; Tye, 2003).

all, it must be understood as relation not between token experiences, nor between contents, but between *phenomenal characters*.

The significance of this point is that a mereological approach like Bayne's is quite appropriate in the pre-phenomenal phase where interdependence occurs and there are detachable parts. Here, grounding interdependence in relations between parts and parts or between parts and wholes is perfectly appropriate. But in the phenomenal phase where phenomenal unity obtains, phenomenal unity is not grounded in a relation between parts, between phenomenal characters, since those relationships were all 'worked out' in the non-phenomenal phase. The phenomenality, and therefore also the phenomenal *unity*, is grounded in the process where that set of interdependent (to some extent) contents becomes phenomenal *en masse*, in Bayne's own words (p. 244).

To clarify further, on the B-first model with which Bayne's view concurs, if we speak of relations (like subsumption) between the phenomenal characters of component experiences, we must make a further distinction. A mental state having a phenomenal character just means a mental state whose contents have acquired the property of phenomenality. We can therefore distinguish two aspects of the concept "phenomenal character". The "character" aspect of "phenomenal character" depends on the content of the mental state, which acquires its particular character, including any interdependence, non-phenomenally. By contrast, the "phenomenal" aspect of "phenomenal character" does not depend on the non-phenomenal content of the mental state, and is only present in the phenomenal phase of the B-first model. Since the character aspect is derived from

the non-phenomenal phase of the model, it cannot be the ground of phenomenal unity. That grounding must inhere in the “phenomenal” aspect of “phenomenal character”. How we think of the nature of this phenomenal aspect does not affect my argument. In order to conclude that Bayne’s holism gives the wrong answer to the parts and grounding questions, it is necessary only for phenomenal unity be grounded in the “phenomenal” aspect of “phenomenal character”, which is a single process of bestowing phenomenality *en masse*, and as such has no parts, and therefore no mereological relations and no subsumption.

The idea that “phenomenality” is what grounds phenomenal unity rather than “character” or relations therewith is further strengthened by that fact that the character of phenomenal experiences can and does change dramatically over time. The thing that remains constant is that bare phenomenality with a particular indexical referent “I” (Kant, 1996, A350).⁴² This is of course just one way of cashing out an NEP view, and there are many others. I offer it as an example of how the NEP position on what grounds phenomenal unity makes more sense of the mechanism of phenomenal unity as revealed by PI—and is less self-contradictory—than Bayne’s holism which suffers for trying to smuggle mereology into holism.

⁴² McMahan’s Theory of Time-Relative Interests is an interesting consequence and perhaps validation of this view. In brief, a person would have little incentive to be egoistically concerned about her future if not for the belief that her present and future selves are identical. While bodies and contents of consciousness change constantly, it is only that bare phenomenality with the referent “I” that remains unchanged. See McMahan (2002) and Schechtman (2014).

3.4. On the parts question

The parts question (Section 2.3) was the point of divergence where NEP holds that there are no parts at all to a total phenomenal experience, while Bayne's holism holds that there are parts but they cannot exist independently of the whole (non-detachable). I argued in Section 2.2 that this kind of metaphysical entity seems unique, and that we should therefore be suspicious of it. It is hard to make sense of Bayne's commitment to marrying mereology with holism without this idea of non-detachable parts. So long as it remains suspicious, the tension persists. If we could find a way to cash out this idea of non-detachable parts, much of that tension would be ameliorated. PI promises to do just that by giving us the idea of an inseparable fusion of parts, an idea I will clarify shortly. This section is devoted to exploring this possible resolution to Bayne's tension, but I conclude that unfortunately it does not succeed, for even inseparable fusion is not the kind of non-detachability Bayne needs to underwrite his marriage of mereology and holism.

A promising way to understand non-detachable parts is the concept of *inseparable fusion*. I argued above that not only bricks, but also quilts are easy to conceive of as having detachable parts that are capable of independent existence. Indeed I have argued that the parts of most kinds of metaphysical entities have detachable parts, with the balance of metaphysical entities not having parts at all (Section 2.2). However, where parts interact deeply with each other, where they melt into each other or even become parts of each other, they are much harder to detach, seemingly impossible to dissect out. Bayne refers to the organs of the human body

coming into being as a whole and therefore needing nothing to bind them together (p. 244). Every medical student discovers just how difficult it is to dissect out from a cadaver the muscles, nerves, blood and lymphatic vessels etc. It takes a great deal of practice, but although it is quite difficult, it is nonetheless also quite possible. So it seems that we cannot ground Bayne's non-detachability just in the phenomenon of fusion alone, for whatever can be fused can always be un-fused.

Another avenue remains for fusion to underwrite non-detachability, this time along the lines of the lone wolf decomposition interpretation (Section 2.4.2). There are situations where detachment is possible only by destroying the part detached. To detach it is for it to no longer be recognisable as itself. The lone wolf is no longer the caribou-killer he was when he was part of the pack. Adding a dimension of fusion to this kind of parthood strengthens the intuition that to remove the part from the whole is to take away its defining properties and thus change its identity. PI means that it is A_c that is subsumed and not A_i , but to detach A from the whole means it is now A_i and not A_c , so it has changed its identity. In this model of inseparable fusion A_c is non-detachable because detaching makes it something else: namely A_i .

But that is not the kind of non-detachability Bayne needs. It could be argued that *every* kind of part is non-detachable in this way, even those we normally think of as quite detachable. Even the brick in the wall loses some of its relational properties when it is removed from the wall: its connections to other bricks, the forces acting upon it, and so on; and so it is no longer exactly the same thing. Yes, the properties lost here are more 'external' to the brick than are the properties lost to A_c , but it

remains the case for both the embedded brick and A_c that they are *not* identical to their detached counterparts. And if that is so, then even this kind of inseparable fusion does not distinguish Bayne's non-detachable parts from the very detachable EP parts. What is more, if we accept that only PPI obtains rather than CPI, another problem arises for Bayne's account. Now we have both interdependent and completely non-interdependent parts of a total phenomenal state to deal with. If Bayne relies on inseparable fusion to underwrite his non-detachability of parts, what is he to say about the non-interdependent and therefore unfused parts? Either he must come up with another way to explain their non-detachability, or he must concede that they are in fact EP-style detachable parts. The latter is devastating for his account, while the former begins to make his account seem somewhat ad hoc—it would surely be better to have a single explanation of non-detachability that covers all parts. So although PI promised to underwrite Bayne's problematic non-detachable parts through inseparable fusion, that approach appears to fail, and the tension in Bayne's account remains.

3.5. Summary

Phenomenal interdependence has illuminated a little further the tension between Bayne's mereology and his holism. It is a principle that Bayne endorses, at least in its limited form if not in its ubiquitous form. It turns out to be potent not only against the phenomenal atomism against which he himself deploys it, but also against his own mereology. On the grounding question, PI highlights the difference between interdependence, which is in essence a property of the non-phenomenal

phase of a token phenomenal experience, and phenomenal unity, which is exclusively a property of the phenomenal phase of a token phenomenal experience. This distinction leads us to ground phenomenal unity not in relations between token experiences, as per Bayne's account, but in the singleness of the event that endows a set of contents with phenomenality *en masse*. On the parts question, PI promised to provide a sound basis for Bayne's problematic non-detachable parts. But on both versions of the idea of inseparable fusion, it did not live up to its promise, and the concept of non-detachable parts remains problematic.

All of this serves to strengthen the case that the tension between Bayne's mereology and his holism remains unresolved. Unless there is some other way of interpreting Bayne, I can only conclude that there is no way to successfully synthesise his mereology and his holism into a single coherent model of phenomenal unity. In the last section below, I consider the implications of this conclusion both for our understanding of phenomenal unity, and of phenomenal consciousness itself.

Conclusion

In this thesis I have argued that there is an unresolvable tension in Bayne's account between his mereology and his holism, and that it is reasonable to conclude from this tension that his attempt to synthesise these two views into a single coherent account of phenomenal unity cannot succeed. I have shown that Bayne's holism diverges from its natural ally NEP on two important points—the parts and grounding questions—and that these two divergences seem to have been moves in the wrong direction. This is borne out by the failure of the verbal interpretation or the decomposition interpretation to make sense of his non-detachable parts, and by the light shone by PI where we saw that the underlying mechanism of phenomenal unity favours the NEP answer to the grounding question, and PI seems incapable of saving non-detachable parts.

What does the failure to resolve this tension mean? If there are no other ways to resolve this tension, then in order to maintain the coherence of his otherwise excellent account, Bayne must sacrifice either his mereology or his holism. Both of these are ominous suggestions. Bayne's mereology, his relation of subsumption (with Chalmers) is an original and brilliant move. It respects some deeply held intuitions about the multiplicity of experience. The genius of his account has been to marry this mereology with another powerful idea that has strong theoretical and empirical support—holistic phenomenal unity. Perhaps there is a way to push this marriage through, but if the arguments in this thesis are sound, and if there is no other way to resolve the tension, then Bayne's account in its current form is not it. Given that NEP views like Tye's one-experience view avoid many of the problems faced by Bayne's account, it seems to me that sacrificing the mereology to preserve an NEP view is the preferable option, but of course that too will have its own problems. Speaking more broadly, the arguments above cast doubt on any approach to phenomenal unity that relies on phenomenal character to individuate experiences (as in Bayne's tripartite conception). This kind of approach has been popular, but as I suggested earlier (Section 2.2), it is an intuition that has arisen out of discussions unrelated to phenomenal unity, and makes assumptions that may be quite safe in other discussions, but which may not be supportable when considering phenomenal unity.

On the bright side, while this thesis has focused on the critique of a tension within the work of Tim Bayne, it has also shown how illuminating the study of phenomenal unity can be for broader questions about phenomenal consciousness generally. Bayne laments that,

“Whether spun from the fabric of philosophical imagination or patched together on the basis of fMRI images, few theories of consciousness are constructed with the unity of consciousness in mind” (p. 225).

Yet it has also been observed that,

“the factors that go into making a particular mental state conscious are inextricably intertwined with those that go into making different state ‘co-conscious’, i.e. go into constituting a unified state of consciousness” (Shoemaker, 2003, p. 58).

Even this brief and inadequate exploration into one small aspect of phenomenal unity suggests some fascinating ways that the study of phenomenal unity can shed light on the very nature of phenomenal consciousness itself.

For example, there is much to be gained from further exploring what more PI can tell us about the mechanisms that produce phenomenal consciousness. If phenomenal unity is more in keeping with a B-first model, this constrains the range of theories of consciousness that are viable, and lends support to models such as *Global Workspace Theories*. It would be interesting to explore how the PI evidence might influence the details of such theories. The idea that contents of experience “become phenomenal *en masse*” would also be fascinating to explore

further in terms of the mechanisms underlying consciousness. These are but some of the questions inspired by the discussion above that are worth pursuing further.

It is amazing how far looking at a blue circle and thinking about it can take you...

Bibliography

- Alter, T. (2010). A Defense Of The Necessary Unity Of Phenomenal Consciousness. *Pacific Philosophical Quarterly*, 91(1), 19–37.
- Anderson, B. L., & Winawer, J. (2005). Image segmentation and lightness perception. *Nature*, 434(3 March), 79–83. doi:10.1038/nature03343.
- Baker, L. R. (2011). Review of The Unity of Consciousness by Tim Bayne. *Times Literary Supplement*, 24.
- Bartels, A., & Zeki, S. (1998). The theory of multistage integration in the visual brain. *Proceedings. Biological Sciences / The Royal Society*, 265, 2327–2332. doi:10.1098/rspb.1998.0579
- Bayne, T. (2000). The Unity of Consciousness: Clarification and Defence. *Australasian Journal of Philosophy*, 78(2), 248–254. doi:10.1080/00048400012349531
- Bayne, T. (2001). Co-consciousness: Review of Barry Dainton's "Stream of Consciousness." *Journal of Consciousness Studies*, 8(3), 79–92. Retrieved from <http://www.ingentaconnect.com.proxy.lib.umich.edu/content/imp/jcs/2000/00000007/F0020001/1142>
- Bayne, T. (2004). Self-consciousness and the Unity of Consciousness. *The Monist*, 87(2), 219–236.
- Bayne, T. (2005). Divided brains and unified phenomenology: a review essay on Michael Tye's *Consciousness and Persons*. *Philosophical Psychology*, 18(4), 495–512. doi:10.1080/09515080500229993
- Bayne, T. (2007). The Unity of Consciousness: A Cartography. In M. Marraffa, M. de Caro, & F. Ferretti (Eds.), *Cartographies of the Mind: Philosophy and Psychology in Intersection*. Dordrecht: (pp. 201–210). Dordrecht: Kluwer Academic Publishers.
- Bayne, T. (2008). The unity of consciousness and the split brain syndrome. *Journal of Philosophy*, 105(6), 277–300.
- Bayne, T. (2010). *The Unity of Consciousness*. Oxford University Press.
- Bayne, T. (2014a). Replies to Commentators. *Analysis*, 74(3), 520–529. doi:10.1093/analys/anu066
- Bayne, T. (2014b). The Multisensory Nature of Perceptual Consciousness. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 15–36). MIT Press.
- Bayne, T. (2014c). The Unity of Consciousness: Summary. *Analysis*, 74(3), 488–490. doi:10.1093/analys/anu054
- Bayne, T., & Chalmers, D. J. (2003). What is the Unity of Consciousness? In A. Cleeremans (Ed.), *The Unity of Consciousness: Binding, Integration and Dissociation*. (pp. 23–58). Oxford University Press.
- Bayne, T., & Montague, M. (Eds.). (2011). *Cognitive Phenomenology*. Oxford University Press.
- Beaney, M. (2015). Analysis. In (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*. Stanford University. Retrieved from <http://plato.stanford.edu/archives/spr2015/entries/analysis/>

- Bennett, D. J., & Hill, C. S. (2014a). A Unity Pluralist Account of the Unity of Experience. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 233–253). MIT Press.
- Bennett, D. J., & Hill, C. S. (Eds.). (2014b). *Sensory Integration and the Unity of Consciousness*. MIT Press.
- Bertelson, P. (1988). Starting from the ventriloquist: The perception of multimodal events. In M. Sabourin, F. Craik, & M. Robert (Eds.), *Advances in Psychological Science, Vol. 2* (pp. 419–439). Psychology Press.
- Bonjour, L. (1998). *In Defense of Pure Reason*. Cambridge University Press.
- Brook, A. (2012). Review of “The Unity of Consciousness”, by Tim Bayne. *Australasian Journal of Philosophy*, 90(3), 599–602. doi:10.1080/00048402.2012.671838
- Brook, A., & Raymont, P. (2014). The Unity of Consciousness. In (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*. Stanford University. Retrieved from <http://plato.stanford.edu/archives/spr2014/entries/consciousness-unity/>
- Brown, J. R., & Fehige, Y. (2014). Thought Experiments. In (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/archives/fall2014/entries/thought-experiment/>
- Calvert, G., Stein, B., & Spence, C. (Eds.). (2004). *The handbook of multisensory processing*. MIT Press.
- Carruthers, G. (2015). Intuitions, edited by Anthony Robert Booth and Darrell P. Rowbottom. *Australasian Journal of Philosophy*, 1–4. doi:10.1080/00048402.2015.1029501
- Connolly, K., French, C., Gray, D. M., & Prettyman, A. (2011). The Unity of Consciousness and Sensory Integration: Conference Report. In *The Unity of Consciousness and Sensory Integration conference*. Brown University.
- Cummins, R. (1998). Reflection on Reflective Equilibrium. In M. R. DePaul & W. M. Ramsey (Eds.), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry* (pp. 113–127). Rowman and Littlefield.
- Dainton, B. (2000). *Stream of Consciousness: unity and continuity in conscious experience*. Routledge.
- Dainton, B. (2007). Coming Together: The Unity of Conscious Experience. In M. Velmans & S. Schneider (Eds.), *The Blackwell Companion to Consciousness* (online.). Blackwell. Retrieved from http://www.blackwellreference.com/subscriber/uid=39/tocnode?id=g9781405120197_chunk_g978140512019717
- Dainton, B. (2014). Unity, Synchrony and Subjects. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 255–285). MIT Press.
- De Gelder, B., & Bertelson, P. (2003). Multisensory integration, perception and ecological validity. *Trends in Cognitive Sciences*, 7(10), 460–467. doi:10.1016/j.tics.2003.08.014
- de Vignemont, F. (2014a). A multimodal conception of bodily awareness. *Mind, Oxford University Press (OUP): Policy F - Oxford Open Option G*, 1–43. doi:10.1093/mind/fzu089
- de Vignemont, F. (2014b). Multimodal Unity and Multimodal Binding. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 125–150). MIT Press.
- Dennett, D. (1991). *Consciousness Explained*. Penguin Books.

- Descartes, R. (1641). *Meditations on First Philosophy*. In E. Haldane & G. T. Ross (Eds.), *The Philosophical Works of Descartes, volume 1*. (1970th ed.). Cambridge University Press.
- Foxe, J. J., & Schroeder, C. E. (2005). The case for feedforward multisensory convergence during early cortical processing. *Neuroreport*, 16(5), 419–423. doi:10.1097/00001756-200504040-00001
- Frith, C. D., & Frith, U. (1999). Interacting Minds--A Biological Basis. *Science*, 286 (5445), 1692–1695. doi:10.1126/science.286.5445.1692
- Gallagher, S., & Zahavi, D. (2012). *The Phenomenological Mind* (2nd ed.). Routledge.
- Gendler, T. S. (2007). Philosophical Thought Experiments, Intuitions, and Cognitive Equilibrium. *Midwest Studies In Philosophy*, 31(1), 68–89. doi:10.1111/j.1475-4975.2007.00154.x
- Ghazanfar, A. A., & Schroeder, C. E. (2006). Is neocortex essentially multisensory? *Trends in Cognitive Sciences*, 10(6), 278–85. doi:10.1016/j.tics.2006.04.008
- Häggqvist, S. (2009). A Model for Thought Experiments. *Canadian Journal of Philosophy*, 39(1), 56–76.
- Hershenov, D. B. (2008). A Hylomorphic Account of Thought Experiments Concerning Personal Identity. *American Catholic Philosophical Quarterly*, 82(3), 481–502.
- Hill, C. S. (1991). *Sensations: A Defence of Type Materialism*. Cambridge University Press.
- Hill, C. S. (2009). *Consciousness*. Cambridge University Press.
- Hill, C. S. (2014). Tim Bayne on the Unity of Consciousness. *Analysis*, 74(3), 499–509. doi:10.1093/analysis/anu053
- Hume, D. (1739). *A Treatise of Human Nature*. (E. C. Mossner, Ed.). 1985 Penguin Books edition.
- Hurley, S. (1998). *Consciousness in Action*. Harvard University Press.
- Irvine, E. (2013). *Consciousness as a Scientific Concept: A Philosophy of Science Perspective*. Springer.
- Jackson, F. (1982). Epiphenomenal Qualia. *The Philosophical Quarterly*, 32(127), 127–136. doi:10.2307/2960077
- James, W. (1909). *A Pluralistic Universe*. The Perfect Library.
- James, W. (1910). *Pragmatism: A New Name for Some Old Ways of Thinking*. The Perfect Library.
- Kant, I. (1996). *Critique of Pure Reason*. Translated by Werner S. Pluhar: Hackett Publishing Company, 1781/1787.
- Koksvik, O. (2014). Three Models of Phenomenal Unity. *Journal of Consciousness Studies*, 21(7-8), 105–131.
- Lee, G. (2014a). Experiences and their Parts. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 287–321). MIT Press.
- Lee, G. (2014b). Extensionalism, Atomism and Continuity. In L. N. Oaklander (Ed.), *Debates in the Metaphysics of Time* (pp. 149–173). Bloomsbury Academic.
- Lee, G. (2014c). Unity and essence in Chalmers' theory of consciousness. *Philosophical Studies*, 167(3), 763–773. doi:10.1007/s11098-013-0196-5
- Lorenz, H. (2009). Ancient Theories of Soul. In (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*.

- Retrieved from <http://plato.stanford.edu/archives/sum2009/entries/ancient-soul/>
- Lowe, E. J. (2002). *A Survey of Metaphysics*. Oxford University Press.
- MacBride, F. (2005). Where are Particulars and Universals? *Dialectica*, 52(3), 203–227.
doi:10.1111/j.1746-8361.1998.tb00050.x
- Marmura, M. (1986). Avicenna's Flying Man in Context. *The Monist*, 69(3), 383–395. Retrieved from <http://www.jstor.org/stable/27902982>
- Masrour, F. (2014a). Unity of Consciousness: Advertisement for a Leibnizian View. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 323–345). MIT Press.
- Masrour, F. (2014b). Unity, Mereology and Connectivity. *Analysis*, 74(3), 509–520.
doi:10.1093/analys/anu064
- McMahan, J. (2002). *The ethics of killing: Problems at the margins of life*. Oxford University Press.
- Menary, R. (2010). Introduction to the special issue on 4E cognition. *Phenomenology and the Cognitive Sciences*, 9(4), 459–463. doi:10.1007/s11097-010-9187-6
- Mendelovici, A. (2013). The Unity of Consciousness. *Philosophical Psychology*, 26(1), 158–162.
doi:10.1080/09515089.2012.663716
- Merleau-Ponty, M. (1962). *The Phenomenology of Perception*. Translated by Colin Smith: Routledge & Kegan Paul.
- Metzinger, T. (2011). The No-Self Alternative. In S. Gallagher (Ed.), *The Oxford Handbook of the Self* (pp. 279–296). Oxford University Press.
- Nagel, T. (1974). What Is It Like to Be a Bat. *Philosophical Review*, 83(4), 435–450.
doi:10.2307/2183914
- O'Brien, G., & Opie, J. (1998). The Disunity of Consciousness. *Australasian Journal of Philosophy*, 76(3), 378–395.
- Overgaard, M. (2006). Introspection in science. *Consciousness and Cognition*, 15(4), 629–633.
doi:10.1016/j.concog.2006.10.004
- Peacocke, C. (2013). Phenomenal Content, Space, and the Subject of Consciousness. *Analysis*, 73(2), 320–329. doi:10.1093/analys/ans125
- Pust, J. (2012). Intuition. In (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*. Stanford University. doi:10.1111/1467-9973.00225
- Raymont, P., & Brook, A. (2009). Unity of Consciousness. In B. McLaughlin, A. Beckermann, & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind* (pp. 565–577). Oxford University Press.
- Robinson, Z., Maley, C. J., & Piccinini, G. (2015). Is Consciousness a Spandrel? *Journal of the American Philosophical Association*, 1(2), 365–383.
- Rodriguez-Pereyra, G. (2014). Nominalism in Metaphysics. In (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*. Stanford University. Retrieved from <http://plato.stanford.edu/archives/sum2014/entries/nominalism-metaphysics/>
- Schechter, E. (2014). Partial Unity of Consciousness: A Preliminary Defense. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 347–373). MIT Press.

- Schechtman, M. (2014). Reflections on Persimials. *The Southern Journal of Philosophy*, 52, 163–170. doi:10.1111/sjp.12070
- Schütz-Bosbach, S., Mancini, B., Aglioti, S. M., & Haggard, P. (2006). Self and other in the human motor system. *Current Biology*, 16(18), 1830–4. doi:10.1016/j.cub.2006.07.048
- Schwitzgebel, E. (2014). The Crazyist Metaphysics of Mind. *Australasian Journal of Philosophy*, 92(4), 665–682. doi:10.1080/00048402.2014.910675
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(03), 417–424. doi:10.1017/S0140525X00005756
- Searle, J. R. (2000). Consciousness. *Annual Review of Neuroscience*, 23(1), 557–578.
- Shimojo, S., & Shams, L. (2001). Sensory modalities are not separate modalities: plasticity and interactions. *Current Opinion in Neurobiology*, 11, 505–509.
- Shoemaker, S. (2003). Consciousness and Co-Consciousness. In A. Cleermans (Ed.), *The Unity of Consciousness: Binding, Integration and Dissociation*. (pp. 59–71). Oxford University Press.
- Smith, N. J. (2012). *Logic: The Laws of Truth*. Princeton University Press.
- Thagard, P. (2005). *Mind: Introduction to Cognitive Science* (2nd ed.). MIT Press.
- Tye, M. (2003). *Consciousness and Persons: Unity and Identity*. MIT Press.
- Van Gulick, R. (2014a). E pluribus unum: Rethinking the Unity of Consciousness. In D. J. Bennett & C. S. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 375–392). MIT Press.
- Van Gulick, R. (2014b). Getting it All Together - Phenomenal Unity and the Self. *Analysis*, 74(3), 491–498. doi:10.1093/analys/anu065
- Varzi, A. (2015). Mereology. In (E. N. Zalta, Ed.) *Stanford Encyclopedia of Philosophy*. Stanford University.
- Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford University Press. doi:10.1093/acprof:oso/9780199933662.001.0001
- Wilkes, K. (1988). *Real People: Personal Identity without Thought Experiments*. Oxford University Press.
- Wilson, D. E., & Reeder, D. M. (2005). *Mammal Species of the World. A Taxonomic and Geographic Reference* (3rd ed.). Johns Hopkins University Press. Retrieved from <http://www.departments.bucknell.edu/biology/resources/msw3/>
- Zahavi, D. (2011). Unity of Consciousness and the Problem of Self. In S. Gallagher (Ed.), *The Oxford Handbook of the Self* (pp. 316–335). Oxford University Press.