The Phenomenal Unity of Perceptual Experience

Stuart Crutchfield

BA(Hons), MLitt

Submitted in fulfilment of the requirements for the Degree of PhD

Philosophy Subject Area

School of Humanities

University of Glasgow

Abstract

This thesis is concerned with phenomenal unity, a salient and important feature of our conscious lives, albeit one which has gone somewhat under-explained. My focus will on the phenomenal unity of perceptual experiences at a time, which is best illustrated by the following kind of example: whilst walking through the park, you notice a particular tree, whose foliage and bark have a particularly intense hue and strange texture. At the same time, there is a strong scent in the air, and you can feel the heat of the sun on your arms. Alongside this, birds are calling loudly… whilst undergoing all these different perceptual experiences, it strikes you that although these experiences correspond to different senses, they seem to be tied together in some important way: there is a unity to your overall perceptual experience at this time.

Even in this mundane case, there are various ways in which my experiences are unified: the various properties of the objects that I represent via vision and touch are unified in the sense that they seem to inhere in the same object (and so are *object unified*), the various objects that I am seeing and touching all seem to me to be located in the same space (and so are *spatially unified*), and further, I can jointly attend to several of these experiences together at will (and so these experiences are *introspectively unified*). Over and above this however, we might think there is a distinct other kind of unity. This is a unity of phenomenology, the subjective character of perceptual experience. There is something it is like for me to hear the birds calling, and there is something it is like for me to see the leaves on the tree. But there is also currently something it is like for me to hear the birds and see the foliage, *together*. Further, this togetherness is such that it is an integral part of my current experience.

This final kind of unity is phenomenal unity, and providing a full description and explanation of this phenomenon will be my task in this thesis. In doing this I will address the following two questions: 'what is phenomenal unity?' and 'how should we explain phenomenal unity?'. I will show that phenomenal unity is best thought of as a relation that holds between token perceptual experiences, answering the first question, and in answering the second question, will consider various reductive and non-reductive explanations of phenomenal unity, before arguing that we should explain phenomenal unity in terms of the unified states being the potential parts of the same overall phenomenal state.

Table of Contents

Acknowledgements

First and foremost, I owe an incredible debt of gratitude to Fiona Macpherson, without whose supervision this thesis would never have taken shape. Many thanks also to David Bain, for many helpful comments, and for reading and commenting on a draft of the thesis, and Tim Bayne, for much helpful discussion of the issues in this thesis. Thanks also to Andy MacGregor for helpful last-minute comments.

The following people have all helped me in some way during the course of my writing this thesis, and I would like to extend my thanks to them all: Albert Atkin, Michael Brady, Carl Brusse, Yuri Cath, David Chalmers, Ben Colburn, Robert Cowan, Tim Crane, Parker Crutchfield, John Cusbert, Barry Dainton, Paul Dimmock, John Donaldson, Douglas Edwards, Nik Effingham, Jonathan Farrell, Akiko Frischhut, Cody Gilmore, Alan Hajek, Benj Hellie, Ross Hetherington, Stephan Leuenberger, Chris Lindsay, Sue Lock, Federico Luzzi, John Matthewson, Mohan Matthen, Neil McDonnell, Kelvin McQueen, Matt Nudds, Casey O'Callaghan, Graham Peebles, Ian Phillips, Dimitris Platchias, Louise Richardson, Jonathan Schaffer, Lizzie Schechter, Susanna Schellenberg, Susanna Siegel, Nico Silins, Barry Smith, Martin Smith, Declan Smithies, Alan Weir, Wayne Wu and Gareth Young.

Thanks also to everyone in the Philosophy Department at the University of Glasgow, which has been an encouraging and supportive place to work for the last few years, and to all in the Department of Philosophy at the Australian National University, which was an incredibly stimulating philosophical environment to work in.

Last, but far from least, I would like to thank my friends and family, most importantly of all my parents, without whom I would never have made it this far.

University of Glasgow

GRADUATE SCHOOL OF ARTS & HUMANITIES

Declaration of Originality Form – Research Degrees

This form **must** be completed and signed and submitted with your thesis.
Please complete the information below (using BLOCK CAPITALS).

Name Stuart Crutchfield
Student Number 0505383c
Title of degree PhD
Title of thesis The Phenomenal Unity of Perceptual Experience

The University's degrees and other academic awards are given in recognition of a student's personal achievement. All work submitted for assessment is accepted on the understanding that it is the student's own effort. **Plagiarism** is defined as the submission or presentation of work, in any form, which is not one's own, without **acknowledgement of the sources.** For further information on what may be considered 'plagiarism', please read carefully the University's Statement at**:**
http://www.gla.ac.uk/services/senateoffice/workingwithstudents/studentdiscipline/plagiarism/universityofglasgowplagiarismstatement/

I confirm that this thesis is my own work and that I have:

Read and understood the University of Glasgow Statement on Plagiarism	☐
Clearly referenced, in both the text and the bibliography or references, **all sources** used in the work	☐
Fully referenced (including page numbers) and used inverted commas for **all text quoted** from books, journals, web etc.	☐
Provided the sources for all tables, figures, data etc. that are not my own work	☐
Not made use of the work of any other student(s) past or present without acknowledgement. This includes any of my own work, that has been previously, or concurrently, submitted for assessment, either at this or any other educational institution.	☐
Not sought or used the services of any professional agencies to produce this work	☐
In addition, I understand that any false claim in respect of this work will result in disciplinary action in accordance with University regulations	☐

DECLARATION:
I am aware of and understand the University's policy on plagiarism and I certify that this thesis is my own work**,** except where indicated by referencing, and that I have followed the good academic practices noted above

Signed……
Date……

"A man's at odds to know his mind cause his mind is aught he has to know it with." – Cormac McCarthy, *Blood Meridian*

Section One – Introductory Chapters

Chapter One – Introduction to The Unity of Consciousness

1.1 Introduction

In our everyday perceptual experience, the following kind of scenario should be familiar to all: whilst walking through the park, you notice a particular tree, whose foliage and bark are of a particularly intense hue and strange texture. At the same time, there is a strong scent in the air, you can feel the heat of the sun on your arms and birds are calling loudly. Whilst undergoing all these different perceptual experiences, it strikes you that although these experiences correspond to different senses, they seem to be tied together in some important way: there is a unity to your overall perceptual experience at this time.

Even in this mundane case, there are a remarkable number of ways in which my experiences are unified: the various properties of the objects that I represent via vision and touch are unified in the sense that they seem to inhere in the same object (and so are *object unified*), the various objects that I am seeing and touching all seem to me to be located in the same space (and so are *spatially unified*), and further, I can jointly attend to several of these experiences together at will (and so these experiences are *introspectively unified*). Over and above this however, we might think there is a another distinct kind of unity. This is a unity of phenomenology, the subjective character of perceptual experience. There is something it is like for me to hear the birds calling, and there is something it is like for me to see the leaves on the tree[1]. But there is also currently something it is like for me to hear the birds and see the foliage, *together*. Further, this togetherness is such that it is an integral part of my current experience.

This final kind of unity is the target of this thesis. I will be endeavouring to discover what this phenomenal unity of our perceptual experiences consists of, and what is responsible for it. I will explore the claim that when had together, our perceptual

[1] The term 'something it is like' has its origins in Thomas Nagel's 1974 paper 'What is it Like to be a Bat?'.

experiences possess some conjoint phenomenology, and whether or not this is more than a matter of these experiences being had by the same subject at the same time.[2]

I will however, be addressing only the issue of the *synchronic* phenomenal unity, i.e. the phenomenal unity of our perceptual states at a time. There is a further question of how our perceptual experiences are unified over time, if indeed they ever are, but that question is outwith the scope of this thesis. From this point on then, I will use the term phenomenal unity to refer to synchronic phenomenal unity, unless otherwise stated.

That this phenomenal unity is a genuine phenomenon, and that it involves anything over and above the two experiences being had at the same time by the same subject, is not universally accepted amongst those working on the unity of consciousness, and thus even positing the existence of this phenomenal unity is a substantive claim requiring supportive. The argument required is also lacking from many extant discussions of phenomenal unity, which simply assume that there is such a thing as phenomenal unity. This issue is something that will be addressed in this first section of the thesis.

As a further caveat to my only examining synchronic phenomenal unity, I will also be limiting myself exclusively to talking about the phenomenal unity of perceptual experiences, as opposed to including other kinds of mental states. Thus properly described, I will be dealing with the *synchronic phenomenal unity of perceptual experiences*.[3] Before proceeding with the question of what phenomenal unity is however, I will go into some necessary background detail on some issues surrounding the study of consciousness. The bulk of this chapter will therefore be concerned with presenting background material, and providing an introduction to the concepts necessary for an examination of the unity of consciousness. This background material will be laid out as follows: in the next section I will lay out some background material on consciousness studies, such as the notions of phenomenal and access consciousness. I will then devote some time to introducing in more detail the various kinds of unity of consciousness mentioned above, such as spatial unity and introspective unity. Following the introduction of these background ideas, I will reintroduce phenomenal unity, in more detail, and also suggest how it relates to the various issues that surround the binding problem. Various concepts employed in the discussion of perceptual experiences will also be of central importance throughout this thesis, and so will be introduced also in

[2] The claim that when phenomenally unified, our experiences have conjoint phenomenology is one made by various philosophers working on the topic of the unity of consciousness, such as Tim Bayne & David Chalmers (2003) and Tim Bayne (2010).

[3] Given that I am talking only about perceptual experiences, I will use 'state' and 'experience' interchangeably, and use one rather than the other for merely stylistic reasons.

this chapter. Also taking a central place in various stages in this thesis will be various experiments conducted on subjects who have undergone what is commonly known as the split-brain procedure. Given these cases will emerge several times throughout the thesis, I will in this first chapter provide an overview of the empirical details surrounding these cases, so that we may move straight into their philosophical significance in later chapters. Finally in this chapter, I will give an overview of the thesis as a whole.

1.2 Consciousness – Some Background Material

Throughout this thesis, I will be dealing with the idea of *phenomenal consciousness*. Conscious states come in a huge variety: perceptual experiences, beliefs, desires, emotional states, imaginings and so on, but given that I am examining phenomenal unity, I am interested in those states which are phenomenally conscious.

Phenomenally conscious states are states for which there is something it is like for the subject to have them. Phenomenally conscious states are described as having *phenomenal character* (often abbreviated to 'character', an abbreviation I will employ also). The phenomenal character of a mental state is the particular phenomenology associated with that state, or what it's like for the subject to have it. Seeing a brown dog will have a particular phenomenal character, as will hearing a bird call, or touching sandpaper.

Paradigm examples of phenomenally conscious states are usually taken to be conscious states such as perceptual experiences, pains, and other bodily sensations, as well as imaginings, and other mental states which are non-perceptual, but which nevertheless have something it is like to undergo them.

There are those who hold that not all conscious states have phenomenal character, and thus not all states are phenomenally conscious.[4] Theorists who restrict phenomenal character in this way typically hold that the kinds of states mentioned above, such as perceptual experiences and bodily sensations, are phenomenally conscious as they have a distinctive or phenomenal character, but thoughts and other cognitive states such as beliefs and desires do not. Those who hold this view allow that states not possessive of phenomenal character are often accompanied by states which do, visual imaginings for example, but conscious thoughts themselves are not. Their opponents, by contrast, are willing to allow beliefs, desires and other cognitive states to be phenomenally conscious,

[4] I am here taking the question of whether a state has phenomenal character to be the same question as whether or not that state is phenomenally conscious. There are those who see these as different questions, but that claim starts a debate which I will refrain from entering.

so that there is something it is like for me to entertain the belief that I am in Glasgow, hope that it stops raining, or desire a cup of coffee. Still others are willing to extend phenomenal consciousness even further, and hold that there is something it is like to realise that the conclusion of an argument follows from the premises.

This is a debate about the reach of phenomenal character, and whether or not a belief can have phenomenal character in the same way as a visual experience can. Those who restrict the domain of phenomenal character to perceptual experiences and other sensory states I will call *conservatives*, and their opponents *liberals*, after Tim Bayne (2009).

> A typical conservative might allow that there is a 'what it's likeness' associated with bodily sensations (aches, pains, orgasms), low-level perceptual states (seeing yellow, tasting sourness, hearing something as approaching), and various affective states (fear, anger, elation), but that's about it as far as 'what it's likeness' extends. A phenomenal liberal, by contrast, might hold that the range of 'what it's likeness' includes not only high-level perception (such as seeing an object as a specific type of car) but also includes such cognitive states as judging that it would be a good idea to go to the south of France in April, wondering whether whales are mammals, and hoping that New Zealand will win at cricket.[5]

This debate between liberals and conservatives has bearing on the unity of consciousness in the following sense: as far as liberals are concerned, it is natural to suppose that there will be a single account of phenomenal unity that applies to all kinds of conscious states, not just perceptual states. Conservatives by contrast, will not expect there to be a unitary account of the unity of consciousness, given that not all conscious states are held by them to be phenomenally conscious states.[6]

I will say a few things here to clarify my own position. Firstly, to reiterate, I will be concerned with phenomenal unity in this thesis, which, we might reasonably suppose applies to all states which are phenomenally conscious. However, I will also be restricting the scope of this thesis insomuch as I will only be concerned with the phenomenal unity of *perceptual experiences*, as I have mentioned above. Given that both sides in the above debate agree that perceptual experiences are examples of phenomenally conscious states, I will not be forced to commit myself to either conservatism or liberalism about the reach of phenomenal character.

---

[5] Bayne (2009) p. 665

[6] This issue is flagged up by Bayne (2010).

Aside from phenomenal consciousness, another notion which will be invoked at various points throughout this thesis will be *access consciousness*.

Unlike phenomenal consciousness, which as we have seen is defined in terms of phenomenal character, access consciousness is a purely functional notion.[7] States which are access conscious are ones where the state or its contents are available for use by the subject in reasoning, action and control of behaviour, and reporting. Of these, reporting is given the least weight by some of those writing on the subject, although perhaps somewhat problematically it may be best practical guide we currently have to the presence of access consciousness.

Next to phenomenal consciousness, access consciousness is a far more perspicuous notion. Some difficulties exist however, in explaining the relation between the two types of consciousness. Of course much of the time our phenomenally conscious states will be access conscious also, hence our ability to reason, report and act on the basis of them. Can the two kinds of consciousness come apart though? This issue may be best illustrated by the question 'can a mental state possess one kind of consciousness without the other?'.

Ned Block considers the two kinds of consciousness to be conceptually distinct, and so holds that it is a conceptual possibility that there could be cases in which a subject has states that are access conscious without being phenomenally conscious, and vice versa.[8] Given the nature of access consciousness there seems a greater likelihood of our discovering cases of access consciousness without phenomenal consciousness, and indeed it is arguable that there are various actual cases which present instances of access consciousness without phenomenal consciousness.[9] Actual cases of phenomenal consciousness without access consciousness however, seem in principle much harder to find, due to our current lack of an independent test for phenomenal consciousness. This issue and its impact on the unity of consciousness will become more salient in Chapter Four where I will discuss it at length in relation to an issue concerning whether or not phenomenal unity can ever break down within a single subject, another issue which will be of central importance in this thesis.

I have explained above that I am going to be concerned primarily with phenomenal consciousness, and will be restricting myself to the phenomenal unity of perceptual states. I have also introduced the notions of phenomenal character, and of access

[7] The paradigmatic definition of access consciousness comes from Ned Block (1995).
[8] Block (1995)
[9] Ibid.

consciousness. I will move on now to give an overview of some key debates concerning perceptual experiences.

1.3 Perceptual Experiences – Some Background

Throughout this thesis, I will be dealing with the phenomenal unity of perceptual experiences, and there are various features to perceptual experiences that merit some background explanation before proceeding with this thesis. It is not my purpose here however to enter the various debates about the various features of perceptual experiences. Rather, I will only briefly cover some issues which have an impact on the issue of the unity of perceptual experiences. I will also point out any assumptions I am making throughout this thesis.

Perceptual experiences are standardly taken to be experiences which *represent* the world in some way, and so part of the contents of a perceptual state will be some *representational content*.

Mental states are representational if they present the subject undergoing them with information about objects, properties or the like. We can say that representational states are states which are 'about' these certain objects or properties. To say that a state has content is standardly shorthand for the claim that it has representational content. Fiona Macpherson describes the idea of states having content in the following way

> A state with content is something that is about, or represents, certain objects, properties or relations and has correctness conditions. The paradigm cases of states with content are propositional attitudes, the contents of which are normally specified in 'that' clauses.[10]

As well as propositional attitudes, representational content can also be ascribed to perceptual experiences. Broadly, the content of a perceptual experience is what is conveyed to the subject by that experience. If for example you are looking at a Cy Twombly painting, you will be presented with an array of diffuse colours and shapes and these can be captured in some proposition which specifies how things seem to you upon looking at the painting. Macpherson goes on to describe how the kind of representational

[10] Macpherson (2006) p. 84

content ascribed to propositional attitude states can be ascribed to perceptual experiences also.

> A notion of content similar to that ascribed to the propositional attitudes can be ascribed to visual experiences in the following manner. When someone has a visual experience it may seem to them that such and such is before them, and the proposition that specifies how things seem also demarcates the content of the experience. Thus, if Morag has a visual experience such that it seems to her that there is a yellow square to the left of an orange circle, then the content of that experience is that there is a yellow square to the left of an orange circle.[11]

So, the contents of the visual experience conveys to the subject that the world has certain features. Of course, the visual experience can be illusory in some way, and then the Cy Twombly painting in the above example may not have these features, but the visual experience nevertheless conveys to the subject that it does. The same goes for hallucination. Macbeth's visual experience as of a dagger before him represented the world as being a certain way, namely with a dagger before him. However, since he was hallucinating, the world was in fact not that way.

There are of course views on which perceptual experiences do not represent the world in this way, and as such do not have representational content. In holding one of these alternative views, one might come to think different things about the unity of consciousness. For example, in holding a naïve realist view of perception, on which perceptual experiences do not represent the world, but have objects in the world as their constituent parts, one might come to question the need for any psychological relation which unifies perceptual experiences, as we would not expect to find any psychological relation which unifies parts of the world.

For my part, I will simply be assuming in this thesis that perceptual experiences do represent the world, and thus have representational content. I will thus not be addressing naïve realist or other views of perception, though the above issue will emerge again in the next chapter, in discussion of Michael Tye's views on phenomenal unity.

As well as taking perceptual experiences to have some representational content, I will be assuming, along with most standard views, that perceptual experiences also have something it is like to have them (in virtue of being phenomenally conscious mental states) and so have some phenomenal character.

---

[11] Macpherson (2006) p. 84

A hotly contested issue regarding perceptual experiences is what the relation between representational content and phenomenal character is. This is a hugely controversial issue, which I will not be attempting to resolve here. I will merely note how adoption of the different positions here will impact on issues of the unity of perceptual experiences.

The available positions in this debate can be marked out broadly as follows: one position states that the representational content and phenomenal character of perceptual experiences are independent, and there can even be perceptual experiences which possess character but no content. In opposition to this view, there are several types of *representationalism*. Representationalism is the thesis that phenomenal properties are determined by representational properties.

One representationalist position is that phenomenal character is in some way determined by content. A stronger position is that character is simply identical to content, and so any description of an experience's content will exhaust its character also. Finally, in opposition to representationalism, though independent of the first position above, is the view that character determines content, the opposition between this final position and representationalism being which has explanatory priority, character or content.

The motivation for the stronger representationalist positions stems from the thought that if two perceptual experiences have the same representational content then they will necessarily have the same phenomenal character. Though this thought is not endorsed by the weaker representationalist, they will still hold that phenomenal character is determined by representational content. This determination is often stated as a supervenience claim: phenomenal properties supervene on the representational properties of an experience, so that sameness of representational properties will mean sameness of phenomenal properties.

This debate between representationalists and their opponents, concerning the relation between content and character, has an important bearing on the unity of perceptual experiences, and so will be flagged up again at certain points throughout the thesis. For instance, the strong brand of representationalism taken by Michael Tye influences his views on phenomenal unity, as we shall see in the next chapter. Also, we shall see that there is a way of describing phenomenal unity (which is the focus of this thesis) on which it makes a difference to a subjects overall perceptual state at a time. This difference may potentially be explained as being a difference in content and/or character,

and so how you see the relation between content and character will influence you here also. This issue will also become clearer in the next chapter, where it will be discussed at length. For now, I will move on to give an introduction to the various kinds of unity that consciousness exhibits.

1.4 Types of Unity of Consciousness

Consciousness exhibits various kinds of unity. Before returning to phenomenal unity, which will be the focus of this thesis, I will introduce and explain several of the other kinds of unity that consciousness exhibits, and which will feature at various points throughout this thesis. These other types of unity will feature in later chapters, for as we will see, some proposals for explaining what is responsible for phenomenal unity involve explaining it in terms of other forms of unity. For now though, I will provide an outline of subject unity; spatial unity and object unity (both of which come under the heading of representational unities) and introspective unity, and will also point out some other kinds of unity that consciousness exhibits which will not be so integral to this thesis, but are worth mentioning. As with my discussion of phenomenal unity, references to any of these other types of unity should be taken as picking out synchronic unity unless otherwise noted.

Subject unity

The first kind of unity that I will give an overview of here is subject unity. My conscious states all instantiate this kind of unity, insofar as they are all states which belong to me, a single subject. States are subject unified then, when they are all had by the same single subject.

What subject unity actually amounts to will depend greatly on what conception of a subject is employed, and what conception of a subject is the correct one, something on which there is little agreement. Subjects (or selves) as conceived of variously as souls, brains, psychological networks, living organisms, or virtual entities.

As far as its importance for my discussion of phenomenal unity, one potential explanation for phenomenal unity is to reduce it to subject unity. The plausibility of this thesis will depend on the plausibility of the view of what a subject is. However, since

there is little if any agreement over what a subject should be considered to be, I will not discuss this possible explanation in much detail.

Representational Unity 1 – Object Unity

Object unity is one form of what Tim Bayne calls 'representational unity'.[12] Representational unities concern not the bearer, or subject of states or experiences, as subject unity does, but the objects or contents (put simply, what the states are about, or represent) of these experiences or states. States are representationally unified to the extent that their contents are integrated with one another. The two forms of representational unity I will be concerned with are object unity and spatial unity.[13]

Object unity concerns the degree to which the contents of a subject's states are integrated around their perceptual objects. For example, my current visual experience as of a coffee cup is object unified insofar as it integrates the contents of consciousness around the perceptual object. This process, by which features are attached to the correct perceptual objects, is known as feature-binding, and occurs not only within sensory modalities, i.e. within vision, but across them also. If I were to pick up the coffee cup, my visual and tactile experiences would be object unified to the extent that they attributed the shape, colour, and texture/feel of the cup all to the same single object.

Representational Unity 2 – Spatial Unity

Representational unities also concern not just the objects of perception, but also the space in which these objects are represented as being located. This second kind of representational unity is known as spatial unity. Spatial unity will be examined in much more detail in Chapter Five, but for now it will suffice to formulate it as follows: two states are spatially unified if and only if the objects of those states are represented as being located in a single common space. Precisely how we should formulate the idea of a single common space will also be addressed in greater detail in Chapter Five.

Introspective Unity

[12] Bayne (2010) p. 10
[13] Ibid.

As well as being subject and representationally unified in various ways, perceptual experiences can also exhibit another kind of unity, concerning their availability for introspection. Again, this notion of introspective unity will be explained much more comprehensively in later chapters of this thesis, but here we can say that two perceptual experiences are introspectively unified iff they are jointly introspectible in a single act of introspection. What a single act of introspection involves will also be spelled out in far greater detail subsequently, but we can say roughly that we are introspectively aware of an experience of ours when we attend to it, usually, though not necessarily with a view to forming a belief or judgement on the basis of it. We do so using the direct, non-inferential awareness we have of our own experiences. An act of introspection will thus involve an instance of exercising this direct non-inferential awareness. The chapters in Section Three of this thesis will examine whether or not either spatial or introspective unity can form the basis for a reductive explanation of phenomenal unity.

Miscellaneous other kinds of unity

As well as being subject unified, spatially unified, introspectively unified and phenomenally unified, our perceptual experiences can also exhibit various other types of unity: they can be *neurophysiologically unified* insofar as they are all realised by the same neural region or via a single neural mechanism.

Experiences can also be *gestalt unified*, where two states are gestalt unified iff a subject's experience of a whole (where this is a whole object or scene) is such that if one salient part of this whole were removed/changed, the subject's experience would have markedly different character. As an example of such a gestalt effect, Tye describes the well known face/vase ambiguous figure shown in Fig. 1.1 On one natural way to see the figure, it is a vase. However, if the right half of the image is removed, then the experience becomes one of a face. Thus, when one salient part of this whole is removed, the subject looking at it will have an experience with markedly different character, that of a face as opposed to a vase.

While interesting, these other types of unity of consciousness will not have a major role in any of the work done in this thesis. The reason for their relative lack of importance here is the fact that they are not plausibly part of any reductive explanations of phenomenal unity, and so they will not be elucidated further in this thesis.

As a final note on the various types of unity described above, though they may apply to various kinds of conscious states, I will be confining myself to discussion of perceptual experiences here also. So unless specifically noted, references to spatial or introspective unity will be to the spatial or introspective unity of perceptual experiences.

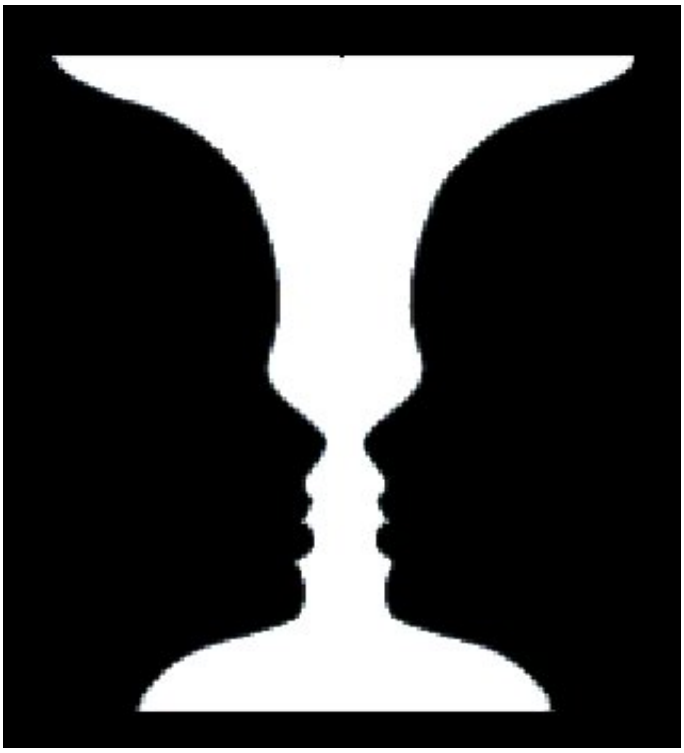

Fig 1.1 Ambiguous figure used to illustrate Gestalt unity.

1.5 Phenomenal Unity

Further to these various types of unity, which have been mentioned above, there is something else, some other kind of unity which our perceptual experiences can possess, and this is the phenomenal unity alluded to in this first paragraph of this chapter, and which is the target of this thesis. Consider again the example we began with above: you are walking through the park, and the colour of the trees, the screech of the birds, the smell in the air and the intense heat all combine together in such a way that your overall perceptual state has some conjoint phenomenology. This conjoint phenomenology that

is present here is as a result of these experiences being phenomenally unified. This phenomenal unity is the subject of this thesis.

This way of introducing phenomenal unity however is less than perspicuous, as it is not immediately obvious what 'conjoint phenomenology' means, and whether it carries any commitments to *extra* phenomenology. Further, there are a wide variety of descriptions of phenomenal unity in the literature on this subject, which reflect a host of different conceptions of what phenomenal unity is. Here is how phenomenal unity is described by Tim Bayne in a (2007) paper.

> Experiences are phenomenally unified when they have conjoint phenomenology; that is, when there is something it is like to experience them together. There is something it is like to have an experience of pain, there is something it is like to see a dog, and there is something it is like to have an experience of a dog and an experience of pain together. One can think of phenomenal unity as a relation that experiences have when they occur as components of a single phenomenal state.[14]

Bayne also describes phenomenal unity in a similar fashion in his (2010).

> Consider again what it's like to hear a rumba playing on the stereo whilst seeing a bartender mix a mojito. These two experiences might be subject unified insofar as they are both yours. They might also be representationally unified, for one might hear the rumba as coming from behind the bartender. But over and above these unities is a deeper and more primitive unity: the fact that these two experiences possess a *conjoint experiential character*. There is something it is like to hear the rumba, there is something it is like to see the bartender work, and there is something it is like to hear the rumba *while* seeing the bartender work. Any description of one's overall state of consciousness that omitted the fact that these experiences are had together as components, parts, or elements of a single conscious state would be incomplete. Let us call this kind of unity – sometimes dubbed 'co-consciousness' – *phenomenal unity*.[15]

The idea of phenomenal unity as involving conjoint phenomenology finds an echo in an earlier paper by Bayne and David Chalmers from 2003.

[14] Bayne (2007) p. 202
[15] Bayne (2010) p. 11

> [T]wo states are phenomenally unified when they have a conjoint phenomenology: a phenomenology of having both states at once that subsumes the phenomenology of the individual states. When A and B are phenomenally conscious states, there is something it is like for a subject to have A, and there is something it is like for a subject to have B. When A and B are phenomenally unified, there is not just something it is like to have each state individually: there is something it is like to have A and B together. And the phenomenology of being in A and B together will carry with it the phenomenology of being in A and the phenomenology of being in B.[16]

All three of these descriptions invoke phenomenal unity's involving conjoint phenomenology, but still this phrase is somewhat ambiguous. The question remains, *'does conjoint phenomenology involve 'extra' phenomenology?'*. When a subject has two or more states which are phenomenally unified, does the conjoint phenomenology of these states also involve some phenomenology that is over and above the phenomenal character of the unified states?

We might call the assumption that concerning phenomenal unity, there is some extra phenomenology that is over and above what it is like to have each of the individual experiences the *extra character intuition*. From the descriptions of phenomenal unity above however, it is not obvious whether Bayne & Chalmers, or Bayne, share this intuition, as it is not obvious whether conjoint phenomenology necessarily means extra phenomenology.

It will emerge in this chapter and the next, that not everyone shares the above extra character intuition, and may thus hold that any problem generated by the question of where and how the extra phenomenology comes from simply dissolves away. In this chapter and the next, I will endeavour to more clearly set up the idea of phenomenal unity, and motivate the thought that there is an issue here that stands in need of further investigation and explanation. Further, I will suggest that even if conjoint phenomenology need not involve extra phenomenology, this is no reason to try and explain away phenomenal unity. As a caveat here however, despite flagging up this issue, I will take no stance on whether or not phenomenal unity necessarily involves extra phenomenology. I will however give some thought to how any putative explanation of phenomenal unity could explain this extra phenomenology were it posited.

Given that we still have ambiguity in the descriptions of phenomenal unity, and given that not everyone shares the extra character intuition, we have potential

[16] Bayne & Chalmers (2003) p. 32

disagreement over precisely what phenomenal unity is. Disagreement that needs at least explained if not resolved, before we can proceed. We need to know what we are trying to explain before we can go about offering explanations of it.

One way to think about what exactly phenomenal unity is, is to consider the following question;

> 'Does the character of the set of perceptual experiences had by a single subject at a time add up to more than the sum of the character of the individual experiences?'

This question addresses the previously mentioned extra character intuition (ECI), which we can re-state as follows:

> (ECI) When two or more perceptual experiences are had at the same time by the same subject, there is something it is like to undergo these experiences which is over and above the mere conjunction of their phenomenal character.

Not everyone shares the extra character intuition, and whether they do or not will influence their answer to the above question. If you think the answer is 'yes', then you are obliged to give some positive account of phenomenal unity, and thus provide an account of this extra phenomenology. Most of the accounts of phenomenal unity I will examine later in this thesis have the means to give an explanation of the extra phenomenology, as I will explain when discussing them.

If you think the answer to the above question is 'no', and so think that phenomenal unity does not involve extra phenomenal character, then you will not be obliged to provide the kind of explanation of phenomenal unity those who endorse the ECI will be. This approach would be attractive to those wishing to explain phenomenal unity simply in terms of subject unity.

Alternatively, answering 'no' to the above question, and rejecting the ECI, paves the way for a move deflationary, or even eliminativist view of phenomenal unity. Such an eliminativist account is given by Michael Tye in his (2003), as we will see in the next chapter.

There are then two questions that can be asked about phenomenal unity:

1. What is phenomenal unity?
2. How should we explain phenomenal unity?[17]

The second question will be the focus of the second and third sections of this thesis. I have started to address the first chapter here, but there is still much more to be said, so I will continue to address this question in the following chapter.

Another important question that I will flag up here, and will return to in the next chapter, is if you hold that phenomenal unity necessarily involves extra phenomenal character, over and above the character of the unified experiences, does it also necessarily involve extra content? If phenomenal unity does mean extra content as well as extra character, what is the nature of this content, and what is responsible for it? How you answer this will depend on how you see the relation between content and character. As such, representationalists will be driven towards the view that if there is extra character involved in phenomenal unity, then there has to be extra content also, given that there would be no other way for the extra character to emerge. If a plausible account of what kind of content cannot be given this in turn might lead the representationalist denying extra character also. By contrast, those who see content and character as being independent of each other may posit extra content or not. This shows that your other commitments in the philosophy of perception may influence your view of phenomenal unity, a point which I will flag up at various other points throughout this thesis, in reference to this and other issues.

I will return to the issue of how best to characterise phenomenal unity in the next chapter. For the remainder of this chapter, I will concentrate on providing more technical background material, and giving an outline of the thesis as a whole.

1.6 The Unity of Consciousness and the Binding Problem

The issue of the unity of consciousness is often taken to be very closely related to what is known as the binding problem. We have seen above, that one kind of unity that consciousness exhibits, object unity, involves features or properties being bound to the same object in perception. The question of how object unity is achieved, i.e. how the redness and roundness of a ball are 'bound' to the same object so that we perceive one

[17] Note that if you do not share the PUI, and thus adopt a deflationary view of phenomenal unity, then your answer to the first question may well be 'nothing'. In which case the answer to the second question may be 'we needn't'.

object with multiple properties, is very similar to how the question behind the binding problem is often stated: given that we know that different properties such as shape and colour are processed in different areas of the brain, what is the mechanism by which these properties are correctly bound to the same object? The difference is that the first question concerns conscious experience, whereas the second concerns brain processing.

This difference goes for other kinds of unity exhibited by consciousness also. The question of how a subject's experiences are spatially unified, so that all the objects represented are represented as being located in a single common space, is a question about the content of the subject's perceptual experiences. Similarly, the question of how a subject's experiences are phenomenally unified is also a question about phenomenally conscious perceptual experiences, not about brain processes.

I am not claiming here that questions about conscious perceptual experiences are unrelated to, or even completely independent from, questions about brain processes. However, answers to binding problems don't give us an answer to the question of how experiences are phenomenally unified. Further, even if an answer to how features are bound together *could* answer questions about object unity, this would not give us an answer to the question of how experiences are phenomenally unified. Experiences which are representationally disunified in some way can still be phenomenally unified, and vice versa, as I will demonstrate in Chapter Five with reference to spatial unity.

1.7 Split-Brain Cases – Some Empirical Background

By way of providing more relevant background material, and as they will re-emerge throughout this thesis, I will in this section provide some background information on the so-called split-brain cases. These cases have become central to discussions of the unity of consciousness, especially to discussions of whether or not phenomenal unity can ever break down within a subject. In order that in future we can move straight to the philosophical import of these cases, I will here outline the relevant history of the split-brain cases, and some of the empirical evidence gleaned from them.

The split-brain operations were developed for and used as a last-resort treatment for epilepsy. The operation involves severing the corpus callosum, and on occasion, the other interhemispheric tracts: the anterior commissure, the hippocampal commissure and the massa intermedia of the thalamus. These interhemispheric tracts are nerve bundles that run between the two hemispheres of the brain, and facilitate high-level

communication between the two hemispheres. When the corpus callosum is severed, the two hemispheres of the brain are seemingly incapable of high-level communication, and most philosophical treatments of the split-brain cases assume that, at least under experimental conditions (which I will outline presently), there is no high-level communication between the two hemispheres.[18]

The first group of these experiments were carried out in the 1930's, but the experimental data comes from trials carried out in the 1950s at Caltech, and in the early 1970s at Dartmouth Medical School.[19]

The importance of the split-brain cases is usually attributed not to how the patients behave in everyday circumstances after the operation, but from the results of putting the subjects in specific experimental situations, where what is sometimes referred to as 'split-brain syndrome' manifests itself. The operation has surprisingly little impact on the patient's everyday life: they can drive, work, and carry out most everyday tasks. This itself may seem surprising, given the disconnection of the hemispheres. The switch between experimental and everyday contexts does have some philosophical importance on some interpretations of these cases, as I will explain in later chapters.

The classic split-brain experiment involves a scenario in which information is presented to the split-brain subject so that it is processed by only one hemisphere. This is usually achieved by means of a tachistoscopic presentation, where the subject focuses centrally, and information is presented for a period of time short enough to preclude the subjects making any eye movement. The subject is often presented with composite words, such as 'key-ring', so that 'key' falls within the subject's left visual field, and 'ring' falls within their right visual field. Since the visual system has a contralateral structure, the information presented to the left visual field will be processed by the right hemisphere, and information presented to the right visual field will be processed by the left hemisphere. Touch has the same contralateral structure, so that tactile stimulus to the left hand side of the body will be processed by the right hemisphere, and vice versa. Smell likewise has contralateral structure, and so under experimental conditions patients may be unable to match odours presented to one nostril with odours presented to the other nostril. Audition on the other hand seems not to exhibit this contralateral structure, information into each ear projecting to both hemispheres. Experiments testing audition

[18] Though there are slight differences in the dissociations that are exhibited by patients with different sections or tracts severed, I will refer to them all as split-brain patients or subjects.
[19] Bogen & Vogel (1962); Wilson et al. (1977); Wilson et al. (1982)

are still carried out on split-brain patients however, and these will be discussed in a later chapter.[20] A typical split-brain experiment is shown in Figure 1.2 below.

The result of an experiment such as that shown in Figure 1.2 below is that the subject will report having seen only the information presented to the right visual field, as this information is processed by the left hemisphere, which dominates control of verbal reporting. They will verbally deny having seen the stimulus presented to the left visual field. At the same time as this, they will use their left hand to correctly retrieve the left field stimulus, the key. When asked to name the object selected by their left hand however, they will respond 'ring'. Although the level of dissociation varies between subjects, this general description holds, and picks out the classic notion of split-brain syndrome.

This provides some background information on the typical split-brain experiments and their results. I will go into more detail about specific cases in later chapters where necessary. The philosophical importance of the split-brain cases comes from their use as putative examples of breakdowns in the unity of consciousness. The non-reductive explanations of phenomenal unity that I will be discussing in the second section of this thesis all hold that phenomenal unity cannot break down, and so all of these views need to offer some account of what happens in these split-brain cases.[21]

I have now given an introduction to the topic of this thesis, the synchronic phenomenal unity of perceptual experiences, and the two questions I will be dealing with: what is phenomenal unity, and what is responsible for it. I have also provided an overview of the relevant background material, and the other debates and issues in the philosophy of mind and perception that impinge on those raised in this thesis. For the final section of this chapter, I will give a brief overview of the chapters to come.

[20] See Milner et al. (1968)

[21] Of course, the claim that phenomenal unity cannot break down is not an essential feature of a non-reductive view of phenomenal unity. It is a claim made however, by the non-reductive views of Bayne, Bayne & Chalmers, and Barry Dainton, which I will be addressing in the second section of this thesis.

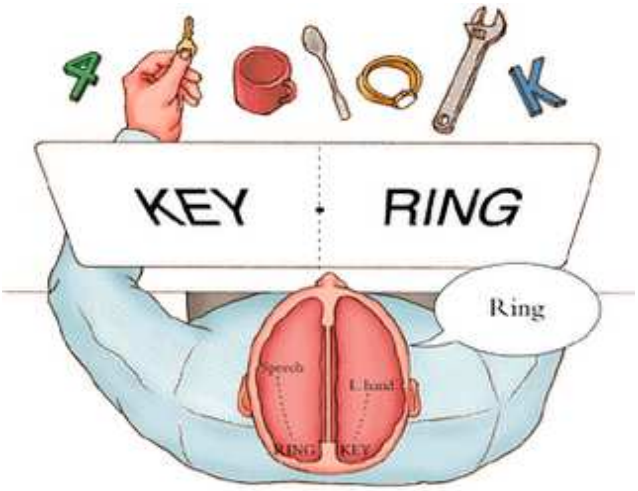

Figure 1.2 This diagram shows a typical split-brain experiment. The subject reports (via the hemisphere controlling verbal reporting) having seen only the stimulus which has been displayed on the right half of the screen ('ring') and further, denies having seen the stimulus presented to the left visual field, and denies recognising the object presented to his left hand. At the same time, the subject uses his left hand to correctly retrieve the object matching the word presented to the left visual field ('key'). When asked to name the object selected by their left hand, he will respond "ring", matching the stimulus to the right visual field.[22]

1.8 Outline of Thesis

In this chapter, I have started on my task of exploring what phenomenal unity is. I have drawn attention to the variety of ways in which consciousness can be unified, and the kind of unity which I am particularly interested in: phenomenal unity. I have outlined what I think is the best way of motivating a debate about what phenomenal unity is, and how we should explain it, as well as suggesting why one might take a deflationary line, and try to explain away the phenomenon. I have also in this chapter given some background on some of the key concepts that will re-emerge throughout this thesis, as well as some background information on split-brain cases, which will also feature heavily throughout the thesis. Before concluding this chapter, I will give an outline of what is to come in the following sections and chapters.

[22] From Sperry (1974)

In the next chapter, I will continue with the task of examining the different conceptions of what phenomenal unity is, and why we should think there is any such phenomenon at all. The focus of this chapter will be the disagreement between Michael Tye and his opponents. Tye's opponents here are all of those who hold that we need to posit some relation that holds between perceptual experiences in order to explain phenomenal unity. Tye claims that this idea is mistaken, and if we analyse phenomenal unity in terms of the closure of content under conjunction, the standard problem of what phenomenal unity is, and how states become so unified disappears. We will see in this chapter the influence of background commitments, as Tye's view is influenced in no small part by his commitment to a strong form of representationalism.

The object of the chapter will be to show that Tye's attempt to explain away phenomenal unity cannot in fact do so, and that he is still obliged to give an account of phenomenal unity. At the close of this chapter then, we will have further explored the issue of what phenomenal unity is, and gone some way toward motivating the claim that some positive account of phenomenal unity is necessary.

Chapter Three will be the first in the second section of the thesis, which deals with the question of how we are to explain phenomenal unity. This section will centre on non-reductive accounts of phenomenal unity. Here I will be taking non-reductive to mean that the views outlined in this section do not seek to reduce phenomenal unity to any other kind of unity that consciousness exhibits. Strictly, only one of the views in this section is non-reductive in the more technical sense that it explains phenomenal unity only in terms of a primitive relation. I will at the opening of this chapter explain in greater detail why I am using non-reductive in these two different ways, and when they both apply. A further component in common between all the views I will consider in this section is the claim that phenomenal unity *cannot break down in a single subject*. as well as an evaluation of how successful primitive relation views are, as an explanation of phenomenal unity, this section will serve as a test of how successful the primitive relation views are in arguing that phenomenal unity cannot break down.

There will be two chapters in this section. The first, Chapter Three, will focus on outlining the particulars of the view that phenomenal unity is a relation that holds between individual token experiences and should not be explained in terms of some other kind of unity.. Such views are espoused by Barry Dainton, Tim Bayne, and Tim

Bayne & David Chalmers and I will devote this chapter to outlining their central claims, and the differences between them.[23]

Chapter Four will also deal with these same non-reductive explanations of phenomenal unity, though this time the focus will be on the problems with these views. As well as the issue of individuating experiences, and a suggestion of some potential explanatory shortcomings with the non-reductive views in this section which explain phenomenal unity in terms of a primitive relation, this chapter will deal with potential counterexamples to the transitivity claims made by Bayne, Dainton and Bayne & Chalmers, and the corresponding transitivity and unity theses. Much of this discussion will centre on the responses made by these theorists to split-brain cases, which may initially seem counterexamples to any claims that phenomenal unity is transitive.

At the conclusion of Chapter Four, I will suggest that none of the three non-reductive views offers a plausible explanation of phenomenal unity, claiming as they all do that phenomenal unity cannot break down. I will claim that none of the three views sufficiently defuse the worries raised by the split-brain cases to hold on to the claims about phenomenal unity not breaking down. A view of phenomenal unity which allows phenomenal disunity has the advantage of not needing to impose one uniform interpretation of any phenomena such as the split-brain cases. Further, with regards to one of the non-reductive views (that of Barry Dainton) that I will discuss in Chapters Three and Four, I will claim that there is something explanatorily unsatisfactory about the positing of a primitive relation. I hope to show that phenomenal unity can be much more fully explained in terms of some other kind of relation holding. Accordingly, such views of phenomenal unity are what I will consider in the third and final section of this thesis, though I will in the end opt to endorse a different approach.

Starting this section, in Chapter Five I will look at possible explanations which reduce phenomenal unity to spatial unity. Initially in this chapter, I will provide some motivation for why such a reductive thesis would have some initial plausibility, pointing out the near ubiquity of spatial content, and, some may think, spatial unity in our perceptual experiences, and why because of this ubiquity, we might think that what is responsible for what is described above as phenomenal unity is simply the fact that our perceptual experiences are spatially unified. This claim forms the basis of the S-thesis, which will be the focus of this chapter

[23] Bayne (2010), Bayne & Chalmers (2003) and Dainton (2006)

Before this kind of thesis can be properly evaluated however, the concept of spatial unity itself needs further unpacking, as there are several different ways in which two states can be spatially unified, corresponding to various ways in which we might think of the notion of a single common space. Thus I will spend some time in this chapter exploring different conceptions of spatial unity, and how they would lead to different versions of the S-thesis.

Unlike the primitive accounts of phenomenal unity, reductive accounts will posit necessary and sufficient conditions for two states being phenomenally unified, and in this case, they will be the very same conditions for two states being spatially unified. The bulk of this chapter then, will be an attempt to show that spatial unity, on any of the ways of conceiving of it, is neither necessary nor sufficient for phenomenal unity. That is, we can imagine cases of a subject's perceptual experiences being spatially unified without their being phenomenally unified and phenomenally unified without their being spatially unified, and I will describe such cases in detail in this chapter. The conclusion of the chapter will be that despite the initial plausibility of the thesis that states that phenomenal unity is reducible to spatial unity, this reduction ultimately fails.

Chapter Six deals with another reductive (in the more general sense) theory of phenomenal unity, which attempts to explain phenomenal unity by showing that it is identical with another form of unity; in this case introspective unity. This will be known as the I-thesis. I will begin this chapter with a caveat: here, and for the rest of this thesis, I am not supposing the truth of any particular theory of introspection, because of a lack of a consensus in philosophy about what the correct theory of introspection is. As with the previous chapter, I will begin this chapter with an outline of why a reductive explanation of phenomenal unity in terms of introspective unity may seem initially plausible.

The particular version of this reductive explanation I will discuss in this chapter is that phenomenal unity can be reduced to joint introspectibility. Two experiences are phenomenally unified then, when they are jointly introspectible in a single act of introspection. This thesis is rejected by Tye, Dainton, and by Cody Gilmore, but I will argue in this chapter that they do not succeed in refuting as they do not consider a version of it that anyone would reasonably endorse, and there are much better versions of a introspective unity thesis.

The problems with Dainton's rejection of this introspective unity thesis are that firstly he fails to do enough to keep separate the notions of *introspected*, and *introspectible*,

and second, his criticisms of the thesis rest too heavily on a disputable notion of the attention-dependence of the phenomenal character of experiences. Gilmore's rejection of the introspective unity thesis centre around an ambiguous use of the term 'introspectible', and furthermore, his versions of the I-thesis are either too restrictive or too liberal, and so would either allow phenomenal unity across subjects, or rule out phenomenal unity in animals and young children, both undesirable consequences. In order to try and meet these desiderata, I will then consider two modified versions of the introspective unity thesis, but despite their advantages over the other theses considered, I will conclude that the commitment to some kind of introspection being involved ensures that no version of the introspective unity thesis can be a successful explanation of phenomenal unity. To conclude this chapter, I will claim that despite there being more plausible explanations of the introspective unity thesis than Tye, Dainton, and Gilmore suggest, this still fails as a potential explanation of phenomenal unity, due to a commitment to introspection of some kind having to be involved.

The seventh and final chapter of this thesis will serve as a defence of an alternative reductive explanation of phenomenal unity. This explanation has as it's the motivation meeting the desiderata seen in the previous chapter, but without invoking introspection. We have seen from the previous chapter what these desiderata for an explanation of phenomenal unity would be: it should not stand or fall alongside some notion of the attention-dependence of phenomenal character; nor should its truth be contingent on the truth of any particular theory of introspection. It should not be so liberal so as to allow two perceptual experiences had by two separate subjects to be phenomenally unified; nor should it be so strict so as to rule out the presence of phenomenal unity in children and (some) animals. This chapter will be an attempt to defend a thesis which meets these desiderata.

The thesis I will defend is: *two (or more) perceptual experiences of the same subject are phenomenally unified if and only if they are potential parts of the same overall phenomenal state.* This will be referred to as the third-state thesis.

This thesis will represent a change of approach, being as it is a explanation of phenomenal unity that does not seek to reduce it to another form of unity, though it is a reductive thesis in the stricter sense, as it is not a view that resorts to a primitive relation to explain phenomenal unity. This will be highlighted, and the different senses of reductive that are at play in this thesis will be flagged up again. It will also emerge that this third-state thesis has much in common with the mereological view of Tim Bayne,

though has the crucial difference that in order to be phenomenally unified states need only be the potential parts of an overall phenomenal state. My third-state explanation will also differ from Bayne's mereological view in that it will not carry any commitments to the transitivity or unity theses.

1.9 Conclusion

In this first chapter, I have done the following: given an introduction to the phenomenon that will be the subject of this thesis, namely the synchronic phenomenal unity of perceptual experience, and suggested that we can distinguish phenomenal unity from various other kinds of unity that consciousness exhibits. I have thus outlined the first task for this thesis, namely to give a better and clearer picture of what we are talking about when we talk about phenomenal unity. The second task is to explain what is responsible for phenomenal unity, and as I have outlined I will be considering various reductive and non-reductive approaches. The rest of this chapter has been taken up with providing some requisite background on concepts that will be employed throughout this thesis, such as phenomenal consciousness, and perceptual content, as well as providing some empirical background to the split-brain cases (which are central test-cases for all of the theories of phenomenal unity discussed in this thesis). Finally, I have given a brief chapter-by-chapter overview of what is to come in the thesis. The task of the next chapter is to delve much more deeply into the issue of what phenomenal unity is.

Chapter Two – Phenomenal Unity, the 'One Experience' View, and the Individuation of Experiences

2.1 Introduction

We have seen from the previous chapter that consciousness exhibits various types or kind of unity and that synchronic phenomenal unity stands in need of some explanation. In the next section of the thesis, I will begin to look at potential explanations of phenomenal unity. Before that however, in this chapter, I will further explore what phenomenal unity is. The main concern of this chapter will centre on a dispute between those who hold that there is a substantive phenomenal unity phenomenon that we are obliged to explain, and those who hold a more sceptical line, and claim that there is no need to posit a phenomenal unity relation. This sceptical position here is taken by Michael Tye in his (2003) book on the unity of consciousness.

Out of this disagreement, I hope to show that in fact even Tye has a question to answer regarding phenomenal unity, and that he too must take a stance on what phenomenal unity is. We will see that given Tye's background commitments to representationalism, he would deny that phenomenal unity could involve any extra content (to the subject's overall phenomenal state), and thus any extra character. Thus, those who are representationalists will most likely deny that the conjoint phenomenology involved in phenomenal unity involves any extra phenomenology. However, given that I will show in this chapter that Tye cannot successfully explain away the need for a phenomenal unity relation, the representationalist will be obliged to give a positive view of phenomenal unity, and may be tempted by a view on which it is simply the co-instantiation of experiences within a subject.

The purpose of this discussion will be to further illuminate what phenomenal unity is taken to be, and to show how other philosophical commitments may affect this. Also, I will show that that even if you deny the extra character intuition outlined in the previous chapter, you still have to give some account of phenomenal unity. Given that it is not clear who in fact does hold that phenomenal unity involves extra character, this is to be expected.

The form of this chapter will be as follows: firstly I will introduce Tye's conception of the 'received view' of phenomenal unity; the way in which Tye takes the received view theorists to set up the problem, and their proposed solution. Then I will look at Tye's criticisms of this received view, and how the received view theorist might respond. I will next introduce Tye's own approach to the phenomenal unity problem, and his 'one experience' view, which is his answer to it. Tye's own view has its problems also, and I will look at these next, suggesting that they are serious enough to jeopardise his position, and that despite much of Tye's own motivation for his view coming from background commitments to strong representationalism and transparency claims about perceptual experience, this is no reason why someone else sympathetic to these views and claims may not reject Tye's one experience view of phenomenal unity. An important issue which emerges in discussion of Tye is that of the individuation of perceptual experiences at a time, and I shall discuss this issue in this chapter. Finally, I will show why, despite his claims, Tye does not succeed in explaining away the need for a phenomenal unity relation.

2.2 The Received View of Phenomenal Unity: Tye's Formulation & Criticisms

We have seen in the previous chapter that phenomenal unity is standardly described as involving conjoint phenomenology, and that experiences are phenomenally unified when they have this conjoint phenomenology. The assumption here is that phenomenal unity is a relation that holds between perceptual experiences.

Michael Tye, in his 2003 book *Consciousness and Persons*, challenges this assumption about phenomenal unity.[1] His own view is that we needn't actually posit any relation that holds between perceptual experiences to explain phenomenal unity. He sets his own view up in opposition to what he calls the *received view* of phenomenal unity. This is not a view of the *solution* to the question of what phenomenal is and how it should be explained, but a view of how these questions standardly arise.

This received view, claims Tye, stems from commitments to there being such things as sense-specific perceptual experiences, and to the claim that in order for there to be phenomenal unity, there must be some relation which is responsible for unifying these sense-specific experiences. We may characterise this received view as follows:

[1] Tye (2003)

> **Received View of Phenomenal Unity**: At a time, a single subject has several sense-specific perceptual experiences. There is something it is like to have these experiences together, and this must be explained by the attribution of some phenomenal unity relation that holds between these separate experiences.

This received view is neutral on what this relation is, but it is the commitment to a phenomenal unity relation that is Tye's target, rather than any particular view of what this relation is.

Phenomenal unity, as it has been conceived by various people, rests on a mistake.[2] Here is how Tye describes what he takes to be the issue the received view theorist highlights.

> Holding a ripe apple in my hand, I experience a red surface and I experience a cold surface. These experiences aren't experienced in isolation however. They are experienced together. This is part of the phenomenology of my experience overall. There is a unity in my experience. In what does this unity consist, given that I am subject to two token experiences, one visual and one tactual?[3]

This view of phenomenal unity comes about, according to Tye, due to the fact that the received view theorist holds that the senses function as largely separate channels of information, and so the problem of phenomenal unity is how we are to explain that though I am undergoing separate sense-specific experiences, they nevertheless possess conjoint phenomenology.

Tye goes on to say that the received view of phenomenal unity holds that if there really is something it is like to undergo all these experiences at the same time, then there must be a maximal, encompassing experience, which includes the other experiences. This maximal experience must be the bearer of the phenomenal unity, and not simply the conjunction of the other experiences. It must be a new experience. As Tye says

---

[2] Tye has as his target here mainly Tim Bayne & David Chalmers (2003) and Barry Dainton (2006) & (2007), though by implication, anyone who holds that a subject has several sense-specific states at a time that stand in need of unification, will be a proponent of the received view according to Tye.
[3] Tye (2003) p. 18

> Seen from this vantage point, the problem of the unity of conscious experience, as it applies to the case of simultaneous perceptual experiences, is, first and foremost, to give an account of the nature of the unifying experience in relation to the other experiences.[4]

By 'bearer of phenomenal unity' it seems that Tye has in mind the conjoint phenomenology that is posited by received view theorists such as Tim Bayne. However, Tye also claims that the maximal state cannot simply be the conjunction of the unified states, but must be an experience in its own right. This may make it sound as though the received view theorist is committed to extra phenomenology, over and above conjoint phenomenology. This however, is not a commitment that a received view theorist has to make, as we saw in the last chapter. This point will emerge again later in this chapter, during discussion of Tye's criticisms of the received view, and their potential responses.

Despite failing to elucidate a potential variation in the received views, Tye continues with claims that such views face several serious objections, the first of which is a problem of regress.

The supposed regress problem is this: if what it is like to undergo the maximal experience involves something in addition to what it is like to undergo the five modality-specific experiences simultaneously, then there must be a unifying relation that is itself experienced. This experience of the unifying relation is not a modality-specific experience, but it is an experience nonetheless, for if it was not, then there would not be anything it is like for the modality-specific experiences to be unified. So, there are six experiences; the five modality specific experiences, and the unity experience. However, as the maximal experience is not simply a conjunction of the other experiences, but another experience in its own right, then there must be an experience that unifies the five modality-specific experiences with the experience of unity, and this starts us on the road to infinite regress.[5] Tye's worry is that there can be no phenomenal unity relation between experiences that explains how the experiences in question are unified, for any such relation would have to be experienced, thus giving rise to the question of how it — the relation of phenomenal unity — is unified with those experiences that it unifies.

Tye outlines another criticism of the received view based around an infinite regress, but a regress of content as opposed to an explanatory regress. The problem here is that if we grant that phenomenal unity is a relation that holds between experiences, as the received view does, and grant also that experiences can have simpler experiences as their

[4] Tye (2003) p. 21
[5] Ibid. p. 22

proper parts, then, says Tye, we will always be faced with the question of what it is that makes the simpler experiences unified with the larger experience, i.e. what makes the parts unified with the whole.

> [C]onsider three simultaneous unified experiences, e1, e2, and e3. If the unity or experienced togetherness of any two experiences requires that there be a unifying experience, then the unity of e1 and e2 requires a further experience E that includes them. Since E is unified with e3, another experience E' is now required. But E' is also unified with E; so a further experience, E", that includes E and E' is needed. And the unity of E" with E and E' necessitates yet another experience; and so on without end.[6]

For every experience we have then, we seem to need a further experience to explain how the first is unified with its parts.

The difference between these two objections may not initially be clear, so I will take some time here to explain the differences.

The first regress objection that Tye outlines relies on the received view theorist holding that conjoint phenomenology involves extra phenomenology. This is not the case, as we have seen, and this point will be raised again in due course. Still, if phenomenal unity makes a difference, Tye claims, then the relation which does the work must itself be the object of experience. If a subject has five perceptual experiences at a time, the relation which unifies them must itself be an experience. But what unifies this sixth experience with the five experiences? There must be a further experience in order to do this. But what unifies this seventh experience with the six experiences? And so on, without end.

The second regress objection concerns not the commitment to extra phenomenology that Tye attributes to the received view theorist, but the coherence of the idea of a maximal unifying experience. For any more two or more experiences, there will always have to be another experience to explain the unity of any two of these experiences. So, the first regress objection of Tye's concerns what he takes to be the received view's commitment to extra phenomenology, and the second regress objection questions the very coherence of the idea of a maximal unifying state.

Tye also puts forward an objection based on an appeal to the supposed transparency of experience. Tye is a champion of the transparency claim and like others, uses it to support a thesis of representationalism about perceptual experience.

---

[6] Tye (2003) p. 22

Tye's representationalist thesis states that phenomenal character is identical to representational content. Transparency is used to support this representationalist thesis, and I will briefly outline how. Firstly, though, the transparency claim itself needs spelling out. Tye elucidates it in the following way:

> Focus your attention on the scene before your eyes and on how things look to you. You see various objects: and you see these objects by seeing their facing surfaces… Intuitively, the surfaces you directly see are publicly observable physical surfaces. They are at varying angles to the line of sight and varying distances away… In seeing these surfaces, you are immediately and directly aware of a whole host of qualities. You may not be able to name or describe these qualities but they look to you to qualify the surfaces. You experience them as being qualities of the surfaces. *None* of the qualities of which you are directly aware in seeing the various surfaces *look* to you to be qualities of your experience. You do not experience any of these qualities *as* qualities of your experience. [7]

When we see objects then, the advocate of the transparency thesis holds that we are not aware of any properties of our experience, only the properties of the objects that we see. Our experiences are transparent in this sense then, we 'see through' them to the objects they represent.[8]

How does transparency affect the received view of phenomenal unity? Tye claims that we are not introspectively aware of our experiences as unified, as we are not introspectively aware of our experiences at all. He takes the received view of phenomenal unity to be undermined by transparency then, as if there are no properties of our experiences given to us in introspection, then the fact that these experiences are phenomenally unified will not be given either, leading Tye to question the basis for thinking that there is such a relation at all.

Note that here there are two positions open to Tye, given his argument from transparency. He could argue that since we are not aware of any properties of our experiences, including their being phenomenally unified, we should not hold that there is a phenomenal unity relation that is given to us phenomenally. Or, he could make the stronger claim, based on the argument from transparency, and hold that since we are not

[7] Tye (2002) p. 138

[8] This transparency claim supports Tye's representationalism in the following way: if introspection does not give us any awareness of any properties belonging to our experiences, rather than the objects they represent, then experiences are made up of represented properties of objects themselves and this must generate the phenomenal character of the experience.

aware of any properties of our experience, including their being phenomenally unified, we should not hold that there is any such unity relation at all. Though this second claim is much stronger, it is this which Tye seems to take from the transparency argument.

> If we are not aware of our experiences via introspection, we are not aware of them as unified. The unity relation is not given to us introspectively as a relation connecting experiences. why, then, suppose that there is such a relation at all?[9]

This certainly seems like the stronger argument outlined above, and a worry is that it may be too strong. It is one thing to deny that a phenomenal unity relation must be given to us introspectively, and another to claim on the basis of its not being given to us introspectively, that it does not exist.

Nevertheless, Tye takes himself to have undermined the received view sufficiently. I will look at the received view theorist's response later in this chapter, but for now I will move on to look at what Tye puts in its place.

2.3 Tye's 'One Experience' View

In place of the received view of phenomenal unity, on which phenomenal unity is some kind of relation which holds between individual experiences in order to unify them, Tye puts his own view of what phenomenal unity is. He refers to it as the *one experience view*. As the name suggests, the main feature of this view is that at a time, a subject does not have several sense-specific experiences that are somehow unified, producing a new, unified experience. Nor are there multiple simultaneous unified experiences within each sense. Thus in looking at the scene before me, and seeing my computer, and a coffee cup on my desk, it is not the case that I am having two or more visual experiences, one as of a computer and one as of a coffee cup.

On Tye's view, there are no such things as purely visual experiences or experiences that are wholly of any one sensory modality. What there is, is a single multi-modal experience, the 'one experience'. This does not so much amount to an answer to the question of how experiences are phenomenally unified, as a sidestepping of it, since on this view of Tye's, there are no separate experiences to be unified, only contents of the one experience. The component parts of the view are the following three claims:

[9] Tye (2003) p. 25

1 – The only experiences that human beings have at a time are entire streams of consciousness.

2 – Experiences cannot have smaller, more basic experiences as their proper parts.

3 – Phenomenal unity should be analysed in terms of the closure of experiential content under conjunction.

Before proceeding, I will give a brief explanation of what it means for content to be closed under conjunction. If we let A and B be experiential contents, then what it means for these contents to be closed under conjunction is for a subject to have an experience which contains A and B. Subject S has an experience as of A, and S has an experience as of B jointly entail that S has an experience as of A and B. I will discuss the potential for analysing phenomenal unity in terms of closure later in this chapter, but for now I will return to Tye's defence of the first two claims above.

The immediate response to this one experience view, Tye concedes, is the incredulous stare. There is, Tye's imaginary interlocutor will say, empirical evidence to support the fact that there are purely visual or auditory experiences, as visual experiences are tokened in the subject's visual cortex, and would still be so if all the other sensory modalities of the subject were malfunctioning.[10] In reply, Tye considers his opponent's position here: in these supposedly purely visual experiences, shape and colour are unified, despite being processed in different areas of the brain. This raises the question "What is the relationship between the experience and its physical basis?".[11] The answer that Tye prefers is that the experience is *constituted* by certain separate and independent physical events that take place in the subject's visual cortex, but the experience is not *identical* with this combination of physical events. As he says:

> On my view, at a given time the wine taster is subject to a single experience that represents the colour of the wine, the sound of the wineglass, as it is flicked, the smell of the wine, and so on. The experience is constituted by a combination of largely independent physical events going on in separate regions of the brain. Within that combination of events, there is a cluster of events (call it C) occurring in the wine taster's visual cortex. In the extraordinary counterfactual situation in which the wine taster's nonvisual senses are all blocked so that no nonvisual information gets in, the wine taster

[10] Appeal to counterfactuals such as this are important when it comes to the issue of individuating experiences, an issue which will be addressed in a later section of this chapter.
[11] Tye (2003) p. 28

> is left with a purely visual experience. And in that counterfactual situation, C, in the absence of the other pertinent actual physical events, constitutes a visual experience. But it does not follow from this that in actual fact C constitutes a purely visual experience. In actual fact, C (wholly) constitutes no experience at all. There is just one unified experience the wine taster undergoes, and C, in conjunction with the relevant events in the auditory cortex, the olfactory cortex, and so on, constitutes that.[12]

Tye considers next the following example, which his one experience view would seem to deal with badly. Suppose you hear a conversation on your left, as you look at a bed of roses, which is in front of you. Intuitively, says Tye's opponent, your auditory experience could have occurred without that visual experience occurring, but on Tye's account, that isn't possible. Tye's reply again draws on the transparency thesis mentioned earlier. When you come to introspect your auditory experience of the conversation, you fail to introspect any *experience* claims Tye. What you actually introspect are the sounds and the auditory qualities that the experience represents. So, the sounds that you experience could have existed without your visual experience of the roses. Further, you could have undergone an experience that represented the sounds, without representing the colours or shapes of the roses. This, claims Tye, is perfectly compatible with his view, on which you only have one experience. In this case, it is simply the represented properties that would vary.[13]

Tye also outlines how his view can cope with experiences in which the auditory component temporally outlasts the visual component. In cases such as this, what happens is that I would initially have an experience which has audiovisual content, and as time passes, it becomes an experience that has only auditory content.[14] Despite it being a feature of Tye's view that the problem of phenomenal unity is sidestepped, Tye is not quite denying that such a thing exists. He is denying that there is a phenomenal unity relation that holds between token experiences, and this is the target of the stronger of the two claims mentioned above. The difference between Tye and the received view is that on Tye's view, phenomenal unity is not a relation that holds between *experiences*; rather it is a relation that holds between qualities *represented* in experience. As Tye says:

> Specifically, phenomenal unity is a matter of simultaneously experienced perceptual qualities entering into the same phenomenal content. The perceptual experience a

[12] Tye (2003) p. 31
[13] Ibid. p. 33
[14] Ibid. p. 34

> normal perceiver undergoes has an enormously rich, multi-modal representational content[15]

What Tye means by 'entering into the same perceptual content' means here is crucial, as this will be what phenomenal unity will amount to on his view.

Further explanation of this from Tye however is not particularly forthcoming. What elucidation there is suggests that two properties' entering into the same phenomenal content stands or falls with the closure of content under conjunction. We can assume then that for two perceptual contents to enter into the same phenomenal content is just for closure under conjunction to hold.

> A consequence of the above position is that phenomenal unity goes with the closure of perceptual experience under conjunction with respect to the unified qualities. Thus, in the case mentioned… in which the loudness of a sound is phenomenally unified for person P with the brightness of a flash of light, the statements 'P has an experience of a loud sound' and 'P has an experience of a bright flash' jointly entail 'P has an experience of a loud sound and a bright flash'.[16]

Phenomenal unity for Tye then, is a matter of closure of content under conjunction. There is no phenomenal unity relation that holds between experiences. There is no phenomenal unity relation that holds between contents either, as given the commitments to representationalism and transparency that Tye holds, contents are just the properties of objects in the world.

Where are we with respect to the first question set out in this thesis then? What is phenomenal unity? According to Tye, the received view says that phenomenal unity is a relation that holds between perceptual experiences and serves to unify them. This view is mistaken according to Tye. In its place he puts a view on which there is no phenomenal unity relation that holds between perceptual experiences, or between experiential contents. Phenomenal unity is simply a matter of a subject having a single experience at a time, the contents of which are closed under conjunction. When closure fails, there we have phenomenal disunity. Tye claims then that he can explain phenomenal unity without positing any relation between experiences or the contents thereof, or positing any extra phenomenology to a subject's overall perceptual state at a time.

---

[15] Tye (2003) p. 36
[16] Ibid. p. 37

For the remainder of this chapter, I will challenge Tye's views. I will not only claim that Tye's one experience view, and his analysis of phenomenal unity in terms of closure is mistaken, but also that even Tye has to give an account not dissimilar from the received view, and that he merely changes the terminology. At the close of the chapter we will see that phenomenal unity is a relation that holds between experiences, and so we will have provided some answer to our first question.

2.4 The Received View's Response to Tye

In the next section, I will present a case against analysing phenomenal unity in terms of the closure of perceptual content under conjunction, but first I will point out some problems with Tye's criticisms of the received view.

Firstly, as I have noted above, Tye's characterisation of the received view is somewhat inaccurate. He seems to attribute to the received view the thesis that phenomenal unity will always involve some extra phenomenology, over and above the phenomenology of the unified states. We have seen however, that the received view theorists, such as Bayne and David Chalmers, do not have to make such a commitment.

Concerning the first infinite regress objection outlined above, the best response that the received view theorist may have to offer is to deny that the phenomenal unity relation is itself an *experienced* relation. There can be something it is like to have a unified set of experiences, without the 'what it is like-ness' including any awareness of the unifying relation itself. The phenomenal unity relation on this line is phenomenal in that it makes a phenomenal difference, but not in the sense that it is itself an object of experience. This is the line of response taken by Dainton, by Bayne & Chalmers, and by Bayne.[17]

In response to this complaint of Tye's, Bayne initially points out that if this objection were to go through, it would count against Tye's one experience view as much as the received view:

> If phenomenal unity is not a relation between experiences, it looks like it must be a relation between the contents of experience, as Tye has claimed… And what would make it the case that simultaneously experienced perceptual qualities – the loudness of a sound, the smoothness of a surface, and the sweetness of a taste – enter into the same

[17] Dainton (2006), Bayne & Chalmers (2003) and Bayne (2010)

> phenomenal content? Arguably, these qualities could enter into a single content only if there were an experiential difference between experiencing the loudness of a sound, the smoothness of a surface, and the sweetness of a taste together as opposed to experiencing these properties separately. But – the objection continues – this surely entails that the subject must be conscious that they are conscious of the sound, the surface, and the taste, and if that's right then we face the task of explaining how that experiential content is unified with the rest of the subject's phenomenology. We appear to have embarked on an apparently vicious regress...[18]

If Tye's objection is successful then, it is too successful, for it affects his view also. This provides us with a hint of what I will conclude in this chapter, that Tye has not avoided the need for a phenomenal unity relation as successfully as he claims.

To return to the received view theorist's response, the way to prevent the regress is to claim, as stated above, that there can be something it is like to have a unified set of perceptual experiences, but without what it's like including any awareness of the unifying relation itself. The inference of Tye's that should be resisted, according to Bayne, is that if there were no awareness or experience of the unity relation itself, then there would be nothing it is like for our experiences at a time to be unified.[19] Why should we not allow that phenomenal unity is a phenomenal relation, insofar as it makes a phenomenal difference, but not in the sense that it has its own phenomenology?

This way of diffusing Tye's worry brings out more clearly the issue surrounding conjoint phenomenology. In Chapter One, I suggested that in denying that conjoint phenomenology necessarily involved extra phenomenology, it would be possible to hold that phenomenal unity consisted simply in the co-instantiation of states in a single subject. Now however, we can see that denying extra phenomenology does not commit you to this line as it is also possible to hold that the phenomenal unity relation makes a phenomenal difference, as it provides the conjoint phenomenology of unified experiences, without this necessarily meaning that the relation has to impart extra phenomenology.

I will remain neutral on whether or not phenomenal unity does involve extra phenomenology as well as conjoint phenomenology, but the explanation of phenomenal unity I will advocate at the close of this thesis will be able to adequately explain any extra phenomenology *if* it is posited, without falling foul of this regress objection of Tye's.

---

[18] Bayne (2010) p. 30
[19] Ibid. p. 31

As for the second regress objection that Tye levels at the received view, here too Bayne has a response that other received view theorists may avail themselves of.

This second regress objection is supposed to undermine the coherence of the idea of a maximal unifying state. Again however, the bite from this objection comes from assumptions that the received view theorist does not have to make. First, as we will see in the next chapter, not all received view theorists need posit a maximal unifying state which subsumes the unified states. Second, as Bayne points out with respect to Tye's argument here, this regress objection rests on the claim that the unity of two perceptual experiences, e_1 and e_m rests on there being a state which is 'bigger' than the two unified states. Bigger here means possessing more phenomenal character. This however is not something the received view theorist need accept. They may hold instead that the unified states are parts of the maximal state, and that since parthood is a reflexive relation, the unifying state is also a part of itself. The state which unifies e_1 and e_m can be e_m itself.[20]

What is the received view theorist to say on the issue of transparency? Of course, there will be some received view theorists who reject transparency, and perhaps also reject the representationalism that transparency is used to support. I do not have the space here to go into the arguments against transparency, so I will not pursue this further.[21]

Bayne suggests however, that even if we grant Tye his transparency claims, and grant that introspection gives us no access to properties of our experiences, only to represented properties of the *objects* of our experiences, we can still maintain that there is a phenomenal unity relation that holds between our experiences. Imagine that you are introspectively aware of hearing a dog barking, and of seeing it run around. Assume that since our awareness of this involves complex phenomenal character, then there must be two experiences here, and they must be unified. The proponent of the received view can still hold here that we have access to the contents of *two unified experiences*. As Bayne puts it:

> If the proponent of the received view were to individuate experiences in terms of phenomenal properties, then she would have a strong case for thinking that any experience with complex phenomenology will involve simpler experiences that are bound together by relations of co-consciousness. In short, the proponent of the received view can respond to

[20] Bayne (2010) p. 29. More detail will be given on the issue of the unified experiences being parts of the maximal state in Chapter Three.
[21] For criticism of Tye's view of phenomenal unity, centring on his representationalism and acceptance of transparency, see O'Dea (2008).

> the argument from introspection even if she is prepared to grant that introspection provides one with direct access only to the representational contents of consciousness.[22]

Bayne has made this point again more recently also:

> [S]ome of the facts that introspection provides access to concern the relations between one's experiences. The reason for this is that one has introspective access to facts about the contents of one's consciousness, and such facts constrain the relations between one's experiences themselves. One doesn't merely have (say) an experience of an apple, a trumpet and an itch; instead one experiences the apple, the trumpet, and the itch 'together' within a single phenomenal field. In order for these contents to be unified the experiences that underlie them must also be unified: no unity in content (or phenomenal character) without unity between the experiences which carry those contents. In short, we can have introspective reasons for thinking that there is a unity relation connecting experiences even if that relation is not directly 'given to us' in introspection.[23]

Of course, Tye would reject this, given his claims about a subject having only a single experience at a time, but that aside, there is at least a way for the received view theorist to reply to Tye's argument from transparency, and without denying transparency.

The received view theorist then, has some reply to all of Tye's objections that I have covered above. Further, as I have shown, there are variations of the received view which are not in fact touched by Tye's objections.

In the final section of this chapter, I will move on to the issue of individuating perceptual experiences at a time, an issue which will turn out to be particularly thorny. Before that however, I will briefly present some considerations given by Bayne against analysing phenomenal unity in terms of the closure of perceptual content under conjunction.

[22] Bayne (2005) p. 11
[23] Bayne (2010) p. 34

2.5 Phenomenal Unity as the Closure of Perceptual Content under Conjunction

Tye's analysis of phenomenal unity is in terms of the closure of perceptual content. He claims that we have no great reason to suppose there is a phenomenal unity relation that holds between experiences (given that a subject has only one experience at a time, which is not divisible into parts), and so insofar as there is phenomenal unity, it is a matter of qualities or properties of experienced objects entering into the same phenomenal content. This in turn is a matter of the closure of perceptual content under conjunction. When a subject's one experience is phenomenally unified, its content will be closed under conjunction, in cases of phenomenal disunity, closure fails.

As a potential problem for Tye's closure analysis, we can consider cases of inter-sensory phenomenal unity, i.e. the phenomenal unity that received view theorists would take to hold between perceptual experiences from different sensory modalities. Bayne formulates the type of case that Tye has to deal with as follows:

> Take a visual experience (V) and a tactile experience (T) that are phenomenally unified with each other. Closure entails that there will be a conscious mental state whose content corresponds to the conjunction of the contents of V and T. Call this state 'P'. what kind of state could P be? Clearly it could not be visual, for were it visual it would not be able to capture T's tactile content; nor, for parallel reasons, could it be tactile. So how are we to think of P? The obvious response is that P is amodal – or, if you like, 'multimodal' – state: it is neither visual nor tactile but visuo-tactile.[24]

So far, this is nothing to trouble Tye. In fact, Tye would point out that since on his view a subject at a time only ever has a single perceptual experience, often with contents from the different sensory modalities, most of the time a subject will have a single multimodal experience. Thus the 'multimodal' response outlined above would be one that Tye would, and indeed does, endorse.

Still, in making this response to inter-sensory unity, Tye may be leaving room to press an objection against closure, and this is the line Bayne initially pushes in response

[24] Bayne (2010) p. 62

to Tye, echoing a similar objection made by John O'Dea.[25] The objection is centred round the following kind of example, which comes from H. P. Grice: imagine that you are holding a coin in the palm of each hand: the coins feel the same size, but they *look* to be different sizes. So, a list of the properties one might be aware of here could be something like the following; the coins as silver, the coins as cool, the coins as being the same size, the coins as being different sizes. The problem here is, as Grice says "…there is nothing in [these] facts to tell us whether the coins *look* different in size but *feel* the same size, or alternatively *feel* different in size but *look* the same size."[26]

O'Dea takes the problem that Grice is highlighting here to be an epistemological problem: the person holding the coins knows that the coins feel the same size, and look different sizes, but there seems no way that they *could* know, based purely on the properties the coins seem to have. There must be then, something other than those properties which carry the information on the basis of which the subject can be aware of which sensory modality is being employed.

Bayne's version of the example is slightly different: consider two scenarios, each involving two coins again, one experienced by you as larger than the other. In the first scenario, you are aware of the large coin by vision, and the small coin by touch. and in the second scenario, the large coin via touch and the small coin via vision. The overall phenomenology of the two scenarios will clearly be different claims Bayne.

> What it is like to enjoy visual awareness of a large coin and tactile awareness of a small coin differs in familiar ways from what it is like to enjoy visual awareness of a small coin and tactile awareness of a large coin. But – so the objection goes – if our account of these two scenarios involves a multimodal representation of the one coin as larger than the other, then we seem to have lost the phenomenal contrast between them.[27]

We might call this the *phenomenal contrast problem* for closure. This kind of thought experiment is also an instance of a more general problem, O'Dea claims. When we both see and touch a single coin, the same question seems to arise as in Grice's case. Since 'circularity' seems to enter twice into the content of our experience, how do we know which is felt, and which is seen?

---

[25] Bayne (2010) and O'Dea (2008)
[26] Grice (1962). p. 136
[27] Bayne (2010) p. 62

As well as being perplexing, the question itself is also peculiar, O'Dea claims. He believes that it is somehow misleading to say simply that 'circular' enters twice into the contents of our experience. When we describe the contents of our perceptual experience, says O'Dea, we leave something out if we list only the properties we are aware of, and ignore the connections between them.

> For example, to describe a visual experience of a red square as simply an experience of an object as red and as square is to miss out something crucial, namely that it is the redness that we are aware of that we are experiencing as square-shaped. It is not the case that we see an *object* which is square *and* which is red – it is the *squareness* which is red and the *redness* which is square. This link is constitutive of the experience itself. In the case of seeing and touching a coin then, although "circular" is in the perceptual experience twice, it is there in two different *ways*: one perceives the object in one's hand as a *silver circle* and as a *cold circle*.[28]

Why then, is this any problem for Tye? It is a problem, O'Dea claims as it means that the question of phenomenal unity is much more tied to the contents of experience than Tye supposes, and perhaps to the contents of sense-specific perceptual experiences, not the multimodal single experience that comes with closure and Tye's one experience view.

On the variation of the coins example that Bayne outlines, there are potential responses. These responses involve taking either a 'Russellian', or a 'Fregean' approach to perceptual content. I will not go into these in great detail here but will limit myself to a brief explanation of how these approaches work.

The Fregean approach involves claiming that even when two different senses represent the same property, size for example, they do so under different modes of presentation, and crucially, via different phenomenal characters.

> The Fregean will deny that tactile and visual experiences of shape (considered as such) have the same phenomenal content, for although the two experiences may represent exactly the same properties they will differ in the ways in which those properties are manifest to the subject.[29]

This Fregean response may not sit well with Tye's representationalism however. Tye holds that character is identical with content, and given that contents are properties of objects in the world, there does not seem room for a difference in content between the

[28] O'Dea (2008) p. 5
[29] Bayne (2010) p. 62

two representations of 'large' or 'small' with respect to the coins. Further, since the different modes of presentation would be modality-specific, this would clash with Tye's one experience view, on which there are no modality-specific perceptual experiences.

The alternative is the Russellian approach, on which there are no modality specific modes of presentation. This approach instead accounts for the difference between the size of the coin as seen and the size of the coin as touched via various representational differences between the senses, for example, texture and temperature via touch, and colour via vision. These representational differences can explain much of the phenomenal contrast between the two scenarios. Further, even if vision and touch do represent the same set of spatial properties, size and shape say, they will do so with varying degrees of specificity or detail. Vision will arguably represent the size and shape of the coin with greater detail than touch, a difference which will be reflected in the subject's experiences. Finally, there may be relational content specific to each sensory modality which will help distinguish a visual representation of a coin of a certain size from a tactile representation of a coin of the same size. Visual experiences may be essentially mediated by movements of the head, and tactile experiences by movements of the hand. This would distinguish a visual from a tactile experience, even if the same property is represented.[30]

So, there are ways that the advocate of closure can respond, but it is not clear that the Fregean response will fit with Tye's brand of representationalism. Instead, Tye may have to endorse the Russellian approach.

This problem of phenomenal contrast is merely the tip of the iceberg claims Bayne, in terms of cases which jeopardise an analysis of phenomenal unity in terms of closure. Closure-based accounts of phenomenal unity will struggle to explain the phenomenal unity of two states of different kinds.

> Suppose that a subject (S) has a conscious judgement with content <p> and a conscious desire with content <q>. Given closure, S must have a conscious state with content <p&q>. But what kind of state could this state be? It couldn't be a judgement, for then it wouldn't capture the fact that the subject *desires* <q>. Nor could it be a desire, for then it couldn't capture the fact that the subject *judges* <p>. But if it is neither a judgement *nor* a desire then it would be unable to capture the fact the subject desires <q> nor the fact

[30] See Bayne (2010) pp. 62-63 for a more detailed description of these two options.

> that it judges <q>. There seems to be no plausible account of the nature of the state with content <p&q>.[31]

I will not pursue this line of objection further however, as it concerns the phenomenal unity of states other than perceptual experiences, and this is outwith the scope I have set for this thesis.[32]

Where are we with respect to closure then? Dealing with perceptual experience only, the most pressing problems come from how we should deal with intermodal unity, and with the phenomenal contrast problem. Given that Tye adopts his one experience view, intermodal unity is not a problem, as a subject will have only a single multimodal experience at a time regardless. As for the phenomenal contrast problem, of the two solutions Bayne identifies, it is not clear that Tye can avail himself of the Fregean approach, and may instead have to opt for the Russellian approach. Again however, what overrules the situation here is Tye's one experience commitment. On the one experience view, rather than there being two pairs of contrasting experiences had by a subject at two different times, there would be a single experience in both scenarios, which at one time would have certain audio-visual content, and at a later time different audio-visual content.

The plausibility of Tye's closure-based account of the synchronic phenomenal unity of perceptual experiences is thus closely tied to his claim that at a time, a subject has only a single multimodal perceptual experience. In the next section I will investigate in more detail the issue of how we should count perceptual experiences at a time, and whether there are any methods for doing so that reflect any joints in nature.

2.6 Individuating Perceptual Experiences at a Time

Phenomenal unity for Tye, as we have seen, simply falls out of the closure of perceptual content under conjunction. Tye's position here also depends in no small part on his one experience claim: a single subject at a time has only one perceptual experience, which is not legitimately divisible into parts.

An initially intuitive way of countering Tye's position would be to attack the one-experience view. However, the basis for criticism of this view begins to look unclear, given the ambiguity in the issue of how we should individuate or count experiences at a

[31] Bayne (2010) p. 64
[32] See Bayne (2010) for more detail.

time. So in this section I will devote some time to an examination of this issue, with the aim of flagging up how this affects the motivation for the phenomenal unity debate, as well as looking at the counting experiences issue itself. Firstly I will take a closer look at potential criteria by which we might individuate perceptual experiences at a time, then I will examine the possibility that any dispute between Tye and his opponents concerning this issue is a merely verbal one, and finally, based on the outcome of testing for a merely verbal dispute, I will outline two ways in which we might continue.

I am currently looking at my computer, and my desk, and a cup. The three objects are spread across my visual field, and are not occluding one another. Do we have here one visual experience, or three, and what settles this? If I am instead looking at a painting of a computer, desk and coffee cup in the same arrangement as the real objects, is this simply a single visual experience, and is this due to there being only one object which I am seeing?

Individuating experiences across sensory modalities may seem a slightly easier task, as we may just say that there are at least as many experiences as there are senses, or something to that effect. However, within a modality, things seem much harder to judge. This issue, namely the synchronic individuation of perceptual experiences, is my focus in this section, and the issue of whether there is any method of answering it that reflects any joints in nature. The intention in this section will be to do the following: flag up that this issue impacts on the more general phenomenal unity question, raise the issue of substantive vs. verbal dispute, and suggest that there may be a plurality of legitimate methods of dividing up experiences (potentially including Tye's method also).

So, to return to the question posed above, how many experiences do I have at time t, whilst seeing the computer, the desk and the cup? Let us suppose that in addition to the visual experience, I am also hearing music coming from the radio, and undergoing tactile experiences caused by the chair I am sitting on. One intuitive way to answer the counting question here may be to say that at time t, I am having as many experiences as I have active sensory modalities: I have one visual experience, one auditory experience, one tactile experience, and so on. On this view, the question of how many experiences a subject has at a time and how we should individuate them has the same answer as to the question of how many sensory modalities there are, and how we should individuate

them.[33] If we adopt this view, the answer to the individuation of the senses question will determine the answer to the 'How many experiences?' question.

Tye, as we have seen, holds a more radical view still, according to which a subject has exactly one perceptual experience at a time. If we want to counter Tye's one experience view, then we need alternative principles of individuation, such that we can demonstrate that these principles give us more than one experience at a time, and crucially, that these are the correct principles, which give us the correct answer.

The following are various potential methods for individuating experiences at a time, that may be used by those who would disagree with Tye:

- Experiences might be individuated according to neural events, with a single neural event realising a single perceptual experience.

- Experiences may be individuated by the instantiation of phenomenal properties, so that an instantiation of phenomenal redness, and phenomenal loudness, would count as two separate experiences.

- Experiences may be divided by appeal to 'phenomenal articulation' to use a phrase of Tim Bayne's (Bayne 2005), where if a subject could have the same experience of the desk that they are having now, without the cup, for example, then these should count as separate experiences. This could also be spelled out by appeal to counterfactuals, i.e., 'if I could have had the same experience of the computer without the bottle, then they are separate experiences'. It seems possible to do this with the other senses: we cannot perceive pitch without volume and timbre in audition, for example.[34]

The question regarding these various potential methods of individuation is whether there is one correct method of individuating experiences at a time which cuts nature at the joints. The worry here concerning the phenomenal unity question is, if there is no one correct way to individuate or count experiences at a time, then we may be able to do so purely according to theoretical need and the influence of background views. This then

33 For more on this question of the individuation of the senses, see Macpherson, F. (forthcoming) "Individuating the Senses", in her (ed.) The Senses: Classical and Contemporary Readings, Oxford University Press.

34 This is by no means an exhaustive list of methods, not is it presented as one.

bars a potential criticism of Tye, as it rules out an appeal to the illegitimacy of his method of individuation.

A further problem is that this question of the synchronic individuation of experiences seems to be such that it is not clear to what we would appeal to get an answer. Appeals to the neural basis for individuation for example can be ducked, as Tye himself does. He in fact denies that there is any neurological basis for the dividing up of a subject's overall perceptual experience, as he denies that perceptual experiences are token identical with neural events and so would deny that there being more than one neural event at a time entails there being more than one perceptual experience.

If it is not clear that there is any way of settling which method of individuation may be the correct one, might this lead us to think that there is not a substantive issue here, and there is merely terminological disagreement? Is it possible that the two sides here (Tye and his opponents) are simply disagreeing over the meaning of the term 'experience'? Though this may initially seem like a line worth pursuing, I will not do so here for the following reason: identification, let alone resolution of verbal disputes is a decidedly tricky affair, and something which is beyond the scope of this thesis.[35]

This all suggests two ways (though not the only two ways) in which we might proceed:

1) If this dispute is a merely verbal one, then Tye's position can be discounted, as he is simply misusing the term 'experience'. This is contingent on our showing that there is a verbal dispute here, and that Tye is in the wrong.
2) If it is not a merely verbal dispute, then we say that there are a multiplicity of legitimate answers to the question of how we individuate experiences, and Tye's position is amongst them.

This first way of proceeding would be one way to show that Tye's position may be rejected, at least inasmuch as we can say that amongst those who are using 'experience' in the standard way, there is still a unity question, as, standardly conceived, a subject has several experiences that need unification. However, given that I have not gone into the issue of identification and resolution of verbal disputes, I will not pursue this first strategy.

[35] For an attempted method for identifying and settling potential verbal disputes, see Chalmers (unpublished manuscript).

At the close of this chapter, I will show that no matter what the result of this debate over the counting of experiences, even Tye still has a unity question to answer. We can even drop all talk of experiences, and the problem remains, for all sides. Before this however, I will take some time to offer a solution to the counting experiences problem, which remains an interesting question in its own right, as well as having relevance to the phenomenal unity question. This is along the lines of the second strategy above.

Supposing that there is substantive dispute between those who hold that a subject has only one experience at a time, and those who hold that a subject has several, can we make any moves to settle this dispute? The second line above suggests that the best way to proceed may be to advance the following view: we are dealing here with a section of reality where there are no privileged joints, and so there are a multiplicity of legitimate individuation criteria for the individuation of experiences. Thus, on this line, there are joints at which we can carve a subject's overall phenomenal state (contra Tye) but, there isn't a unique set of joints.[36] This position would have some benefits: we intuitively think that experiences in different modalities can be unified in the same way that experiences in the same modality can be, and so if there is a legitimate set of joints here, then we can explain this intuition. On this view, we can also posit limits to the joints: though we can legitimately divide up experiences to correspond to the different sense modalities, there are certain divisions that we cannot make, due to facts about perceptual experiences. We cannot, for example, divide up experiences so as to have a visual experience that represents shape without colour. There would seem to be other likely limits, corresponding to the other senses. For example, we could not divide things so as to have an auditory experience that represented pitch without volume or timbre, or indeed any one of these three properties without the other two. More tentatively, we may not be able to divide up experiences so as to give us a tactile experience as of texture without pressure. These limits will correspond to what we might call the laws of perception, and though the full extent of these laws has not yet been worked out, it is reasonable to suppose that philosophical and empirical work can one day accomplish this. For the moment, we can say that the limits on individuation are set by the laws of perception, whatever they may be.

This line is I believe, the one to hold on this issue at present: though there are no currently agreed upon unique set of joints at which to carve, so as to settle the question of how many experiences a subject has at a time, there are limits to the carving: we can

[36] Thanks to Susanna Siegel for this point.

currently see that there are some ways to carve up experiences, that both sides can agree on as being illegitimate, with these limits stemming from features of the format of the experiences that the various sensory modalities present us with.[37]

Though this provides an answer to the question of how we should count perceptual experiences within a subject at a time, it does not provide us with much to put pressure on Tye, for if there are a multiplicity of different but legitimate ways in which we can count experiences, then why is Tye's position not counted as one of these ways? Of course, if his position is not one of the legitimate ones, then we have a reason to discount it, but we have no great reason to discount Tye's view on these grounds. If, then, as it seems, Tye's position is to be included in the group of legitimate positions, then perhaps independent concerns will have to be brought to bear on Tye, if we are to discount his stated position on phenomenal unity, on which there is no question to be addressed. The dialectical use of the view that there is a multiplicity of correct methods for the individuation of experiences will be examined in both the following chapter, and in subsequent discussion of the split-brain cases.

2.7 Conclusion

Where does this leave us then? We have seen that much of the plausibility of Tye's conception of phenomenal unity depends on the plausibility of his claims about a subject having only a single perceptual experience at a time. We have seen above however that assessing the plausibility of Tye's way of individuating experiences at a time will not be straightforward, and in fact it may be best to treat things as though there are several equally legitimate methods of individuation. We have seen that there may be reasons to be doubtful of an account of phenomenal unity based on the closure of perceptual content under conjunction, and we have also seen that Tye's criticisms of the received view of phenomenal unity are not conclusive. My aim in this chapter was to further illuminate the question of what phenomenal unity is, via discussion of Tye and the received view, and to show that even Tye still has to answer the unity question in a way close to the received view theorist.

The reason that Tye still has the unity question to answer is explained by the following: if we take a split-brain patient under experimental conditions, and a non-split-brain subject under normal conditions, what is the mental difference between these two

[37] This issue will arise again in Chapter Three, where Tim Bayne advances a similar claim to mine.

subjects? Of course, there is the obvious physiological difference, but supposing we rejoined the corpus callosum, this would make some mental difference, which would be responsible for the split-brain subject's experiences being phenomenally unified. Tye still has to give an account of what this mental difference is. In other words, what does this physiological difference do? If the difference it makes is that it means that separate contents are contained within the same one maximal experience, then this is the phenomenal unity issue. These contents are still ones about which certain crucial counterfactuals are true, that they could have existed independent of various other contents contained in the one experience, and that they interact with belief in the appropriate way (and these Tye would agree with also). So, regardless of what you call them, there is still a unity question. Tye's arguments discussed in this chapter show only why he objects to the standard terminology, not why there is no phenomenal unity issue.

Yet another way to frame the issue, so as to put the pressure back on Tye vis a vis the phenomenal unity issue, is to point out that Tye is simply shifting the focus away from individual mental states, or pairs/sets of mental states, and onto subjects. In effect, Tye replaces the question 'what makes two states phenomenally unified?' with 'what makes two states of the same subject?'. However, this is in effect the same question.

So, for these reasons, I will be proceeding in this thesis under the assumption that there is a substantive issue as to what phenomenal unity is, but that we should all hold that it is some kind of relation that holds between perceptual experiences, and which is responsible for the conjoint phenomenology of these experiences. I have not however, made a commitment as to whether or not conjoint phenomenology necessarily involves extra phenomenology, and nor will I in the rest of this thesis, though I will ask of any potential explanation of phenomenal unity if it could explain extra phenomenology. As to what kind of relation phenomenal unity is, that will be the concern of the next two sections of this thesis. The final point in this chapter is to note that I will be sticking from now on to the received view terminology, and talking of a subject as having multiple experiences at a time, and taking it that a subject's overall experience can be legitimately divided into parts. This itself will become very important at a later stage.

Section Two – Non-Reductive Accounts of Phenomenal Unity

Chapter Three – Non-Reductive Explanations of Phenomenal Unity

3.1 Introduction

In the previous section of this thesis, my concern was to give some introduction to the notion of phenomenal unity of perceptual experiences, and the various conceptions of it. I concluded the section by claiming that the best way to think of phenomenal unity is as some kind of relation which holds between experiences, and which is responsible for the conjoint phenomenology of these experiences. In this section and the next, I will shift from *descriptions* of phenomenal unity to potential *explanations* of it.

I will consider two broad kinds of explanation of phenomenal unity, *reductive accounts*, where phenomenal unity is explained solely in terms of some other psychological feature or other kind of unity exhibited by consciousness, and *non-reductive accounts*, which explain phenomenal unity without reducing it to any other feature, psychological or otherwise.

Reductive explanations will be covered in Section Three of the thesis, but first in this section I will explain non-reductive views of phenomenal unity. In doing so I will be drawing on the views of Barry Dainton, Tim Bayne & David Chalmers, and on Tim Bayne's own view, developed more recently.[1] All of these views have much in common, hence my treatment of them here under the same heading. Covering all of these non-reductive views here will avoid repetition, and also help better map out the philosophical space. Crucially, all of these views treat phenomenal unity as a relation which is *transitive*, and hold that this transitivity cannot fail. This transitivity claim, if true, in effect means that phenomenal unity cannot break down within a single subject. It will emerge however, that this claim is decidedly controversial. A further commonality between all of the non-reductive views is the requirement for an answer to the problems generated by split-brain cases. Again, to avoid repetition and to better give an idea of the philosophical

[1] Dainton (2006), Bayne & Chalmers (2003) and Bayne (2010)

lay of the land, I will consider the various interpretations of the split-brain cases in the next chapter, after first outlining the various non-reductive theories. Note here that although the three views mentioned above all hold that phenomenal unity is a transitive relation, it is not the case that a commitment to transitivity is entailed merely be posing a non-reductive view of phenomenal unity. As a further note, it should be mentioned here that although I am grouping all these views in this chapter together as non-reductive views, this is for dialectical purposes, as there is a sense in which the views of Bayne and Bayne & Chalmers are actually reductive views. The point to bear in mind here is that as these views do not seek to reduce phenomenal unity to another kind of unity that consciousness exhibits, I am referring to them here as non-reductive. I will however flag up at various points throughout the remainder of this thesis the sense in which Bayne and Bayne and Chalmers are actually proposing reductive explanations of phenomenal unity.

The form of this chapter will be as follows: firstly I will draw out a potential point of difference between non-reductive views of phenomenal unity, whether we should think of think of the phenomenal unity relation as a *'top-down'* relation, or as a *'bottom-up'* relation, who holds these different positions, and what the resultant differences created by these options are. Second, I will outline what it means to say that phenomenal unity is a subsumption relation, and why this prompts Tim Bayne to describe his account of phenomenal unity as a *mereological* account. Here I will also point out, as noted above, that whilst views which analyse phenomenal unity in terms of subsumption are non-reductive in the sense that they do not attempt to reduce phenomenal unity to spatial unity, they differ importantly from the view that simply defines phenomenal unity as a primitive, sui generis relation. I will then outline some of the properties attributed to the phenomenal unity relation by these non-reductive accounts, namely *reflexivity, symmetry and transitivity*. Of these properties, transitivity will prove to be the most controversial, and the hardest to sustain claims of phenomenal unity possessing this property. As mentioned, holding that phenomenal unity is a transitive relation has the consequence that phenomenal unity cannot break down within a subject. This leads some proponents of a non-reductive view of phenomenal unity to advance theses to the effect that this will always be the case, for any two phenomenal states of a single subject. I will in this chapter examine these theses.

Finally in this chapter I will look at the issue of where non-reductive views of phenomenal unity stand on what phenomenal unity is and some of the features it is taken

to possess: are non-reductive views committed one way or another on the issue of conjoint or extra phenomenology.

Though all of the views considered in this chapter cover non-reductive views of phenomenal unity, inasmuch as they do not reduce phenomenal unity to some other kind of unity, or hold that explanations of phenomenal unity would simply fall out of an explanation of feature-binding, for example, there are differences between the potential non-reductive accounts.[2] These differences concern issues such as what properties are attributed to the phenomenal unity relation, and whether phenomenal unity is a *'top-down'* or a *'bottom-up'* relation.

3.2 Top-Down or Bottom-Up?

All of the non-reductive views addressed here take phenomenal unity to be a relation that holds between individual perceptual experiences, contra Tye. These individual experiences are joined by the phenomenal unity relation into a unified whole, the subject's overall or maximal experience at a time. The most immediately apparent difference between top-down and bottom-up views of phenomenal unity is which of these things they take as their starting point, the individual experiences or the maximal experience. The bottom-up view is held by Barry Dainton, and the top-down view held by Tim Bayne, and by Tim Bayne & David Chalmers. Bayne says the following to highlight the difference,

> Whereas treating phenomenal unity as a primitive provides us with a 'bottom-up' approach to the unity of consciousness, one that starts with the multiplicity in consciousness, taking subsumption as our primitive is to adopt a 'top-down' approach to the unity of consciousness and begin with the unity that subsumes this multiplicity.[3]

I will consider the issue of taking phenomenal unity, or subsumption to be the primitive later in this chapter, but here I will say a little more about top-down and bottom-up approaches. In the passage quoted above, Bayne talks of the two approaches starting from different points, and this may be the most obvious difference between the two approaches, what they take to be the explanatory starting point, and thus what has

[2] As mentioned above, I will later in this chapter discuss the difference between views which claim that phenomenal unity is a primitive, and those which analyse it in terms of subsumption, and will suggest that they are both non-reductive in the sense I mean in this thesis.
[3] Bayne (2010) p. 20

explanatory priority. The top-down approach takes as its starting point the subject's maximal experience at a time, and thus asks what relation this maximal state bears to the individual states which it unifies. The bottom-up view on the other hand takes as its starting point the individual perceptual experiences had by the subject at a time, and thus asks what relation these individual states bear to the maximal state, which unifies them.

Given how the two views are described by both sides, there doesn't seem any reason to think that either view is taking a stance on what is *metaphysically prior*, only what the *explanatory* starting point is. Contrast this with the top-down view of Tye, who *does* hold that the maximal state has metaphysical or ontological priority, claiming as he does that the maximal state cannot be divided. I will not be discussing Tye's view further here however, as I have addressed it at length in the previous chapter.

Further evidence that the difference between top-down and bottom-up approaches is merely in the explanatory starting point comes from Dainton's claims that the two views are not incompatible.

Dainton discusses the relation between the two views in his (2007) paper: he points out that Bayne and Chalmers' view has its benefits from the point of view of formalisation; their view can be evaluated in terms of more familiar logical relations such as entailment, which offers more by way of explanation, whereas his own view has to remain with the primitive relation that holds between individual states.[4] But, Dainton claims, the two views are not in competition, and are certainly not inconsistent with each other. He points out that any set of states that are jointly related by the phenomenal unity relation that he posits will form a whole which subsumes its parts, so the two views tacitly involve each other.

> To appreciate this it suffices to pose the question "If a state S subsumes states S_1, S_2… S_N, just what is S itself like?" In spelling out the answer we will mention the qualitative character of the various constituents subsumed in S; if some of these constituents are spatially related or introspected we will mention this too. Is this enough? Arguably not, for we have not yet made explicit the manner in which the states subsumed in S are unified… for our description to be phenomenologically adequate – for it to capture what it's like to have S – it will also need to mention that each part of S is connected to every other part by the relationship of experienced togetherness. Hence not only does co-consciousness* remain very much present (albeit tacitly) in the subsumption approach, it

[4] Dainton (2007) p. 217

> is (arguably) responsible for conscious states having the distinctive form of unity that they do.[5]

So, the top-down and bottom-up approaches to unity are not incompatible according to Dainton. I will proceed then under this assumption, that the top-down and the bottom up views of phenomenal unity are not incompatible, and the different starting points of the two views reflect only explanatory differences, not metaphysical ones.

3.3 Primitive Relations, Subsumption and Mereology

As mentioned above, there are non-reductive views which take phenomenal unity to be a primitive, such as Dainton's, and there are non-reductive views which give an analysis of phenomenal unity in terms of subsumption. I will return to the question of whether these should both be classed as non-reductive views shortly, but first we need more information about what it is to say that phenomenal unity is a primitive relation, or that it is a subsumption relation.

Treating phenomenal unity as a primitive is the tactic taken by Dainton. He explains at length his reasons for rejecting reductive views of phenomenal unity, and instead advocates a primitive relation view.[6]

> [I]f experiences can be phenomenally unified without being unified in any other discernible way, are we not forced to recognise the existence of a "pure" relationship of co-consciousness?[7]

The phenomenal unity relation is, according to Dainton, a basic feature of perceptual experience. It has no phenomenal features of its own, nor is it due to some separate ingredient which is added to perceptual experience.

Treating phenomenal unity as a primitive relation, there is not much more to say. Phenomenal unity is a feature of perceptual experience, it is a relation which holds between individual experiences, and it has certain logical properties, which I will examine

[5] Ibid. (Note: Dainton uses 'co-consciousness' synonymously with 'phenomenal unity'.)
[6] These reasons will be examined in the next section of this thesis.
[7] Dainton (2007) p. 215

shortly. Is this enough, for an explanation of phenomenal unity? I will return to this question in due course.

An alternative view, which allows us to say more by way of explanation, is to hold that phenomenal unity is a relation of subsumption. This approach is taken by Tim Bayne, both in his (2003) paper with David Chalmers, and more recently in his (2010) book.

In Bayne's own account subsumption is explained in terms of parts and wholes. The central claim of Bayne's account is the following:

> [T]wo conscious states are phenomenally unified when, and only when, they are co-subsumed.[8]

This central claim needs unpacking. Firstly, what does Bayne mean by subsumed and subsumption here? As mentioned above, Bayne thinks of subsumption in terms of parts and wholes, that is, in mereological terms. So what it is for one experience to subsume another is for it to contain the subsumed experience as a proper part.

This leads Bayne to claim that a subject's total experiential state at a time is a whole, which includes as parts various experiences, such as the olfactory experience of smelling coffee, or the visual experience of seeing a bird. So, one's overall experience contains within it other, 'smaller' experiences. From this, it is easy to see the relation between subsumption and holding a top-down view of phenomenal unity.

This is all Bayne has to say on the notion of subsumption here, and with it the claim that the phenomenal unity relation is one of subsumption. More elucidation of subsumption can be found in Bayne & Chalmers' (2003) paper.

Bayne & Chalmers, in elucidating their notion of subsumption, advance a position which initially seems very similar to Dainton's. They initially posit the notion of subsumption as being something like an intuitive primitive. A paradigm case of subsumption on this view is taken to be the relation between a complex phenomenal state, and a simpler state that might intuitively be seen as one of the more complex state's parts.[9] They also stipulate that subsumption holds between a phenomenal state and less specific states that correspond to the same perceptual experience, so that seeing a blue ball subsumes the state of seeing a ball. This sort of subsumption is necessary they claim,

[8] Bayne (2010) p. 20

[9] See Chapter Two for an overview of the issues surrounding the carving of phenomenal states such as experiences into parts, and Tye's opposition to this.

for it to be possible for a highly specific overall phenomenal state to subsume all of a subject's current phenomenal states, including unspecific ones.[10]

Bayne & Chalmers do point out that alternative analyses of phenomenal unity as a primitive relation are available, and cite Dainton's view here. They claim that their view of phenomenal in terms of subsumption "runs deeper" than an analysis in terms of a primitive relation holding between experiences, presumably as they are able to appeal to subsumption in their explanation, as opposed to a primitive relation. I will later in this chapter return to discussing the differences between the two respective positions.

The notion of subsumption in Bayne & Chalmers' position is connected to the notion of 'what it is like' in the following way, they claim: when A subsumes B, what it is like to be in B is an aspect of what it is like to be in A. an alternative explanation of subsumption in terms of 'what it is like' would go as follows:

> *A phenomenal state A subsumes phenomenal state B when what it is like to have A and B simultaneously is the same as what it is like to have A.*[11]

There is a close relation, claim Bayne and Chalmers, between the notions of subsumption and entailment. They formulate entailment as follows: for two states P and Q, P entails Q when it is impossible (either logically or metaphysically) for a subject to instantiate P without instantiating Q. So, if P and Q are phenomenal states, when P *subsumes* Q, then P will *entail* Q also. Bayne and Chalmers give the following as an example:

> For example, if P involves the phenomenal character as of seeing a red book and hearing a bird singing, and if Q involves the phenomenal character as of seeing a red book, then it is impossible to have P without having Q.[12]

As it goes for this example, so it goes for all cases of subsumption claim Bayne & Chalmers: the subsuming states always brings with it the subsumed state.[13]

Despite the closeness between entailment and their notion of subsumption however, Bayne and Chalmers do not define subsumption in terms of entailment. So although phenomenal unity is to be analysed in terms of subsumption, subsumption here is not simply entailment. This discussion suggests some of the differences between the

[10] Bayne & Chalmers (2003) p. 40
[11] Bayne & Chalmers (2003) p. 41
[12] Ibid.
[13] Ibid.

account of Bayne and Chalmers, and that of Dainton: despite not wishing to reduce phenomenal unity to some other kind of unity, Bayne and Chalmers (and Bayne himself) do provide some analysis of the phenomenal unity relation, as we have seen. Dainton, on the other hand eschews further analysis of phenomenal unity in favour of characterising it as a primitive, sui generis relation. There is then, a question about whether or not these views are non-reductive views, and I shall return to this point in due course.

As I have covered what Bayne and Chalmers jointly say about subsumption, I will leave the issue here, and focus on some other points concerning Bayne's own account, assuming that the two positions are equivalent on this point, and Bayne means the same by subsumption in his more recent work as he does in his work with Chalmers. As an aside here, though he has changed his interpretation of the split-brain cases, as we will see, the technical details of Bayne's analysis of phenomenal unity as a subsumption relation are as they were in the Bayne & Chalmers paper, and as such I will treat the views as more or less the same, unless explicitly stated, but will refer for the most part to Bayne's more recent work, again, unless stated.

We have a better idea then, of what it would mean to claim that phenomenal unity is a primitive relation, or a relation of subsumption. However, there is still the question which I raised earlier in this chapter, are both of these truly non-reductive views of phenomenal unity?

In the case of Dainton's primitive view, the answer seems quite clearly 'yes'. On Dainton's view phenomenal unity is a basic feature of perceptual experience, a sui generis relation about which nothing further can be said in terms of explanation. In the case of the subsumption view, things are trickier to judge. Dainton himself seems to see the view of Bayne & Chalmers as treating phenomenal unity as a primitive feature of experience, but there may be reason to think otherwise.[14]

The subsumption view does not claim that phenomenal unity is identical with some other form of unity exhibited by consciousness, so in this sense it is a non-reductive view.[15] However, it might be argued that you simply "get phenomenal unity for free" out of an analysis of subsumption (as outlined above). On this view, phenomenal unity would simply fall out of the subsumption relation which holds between parts and wholes of the maximal perceptual experience had by a subject at a time. This would be in much the same spirit as Tye's view that phenomenal unity simply falls out of the closure of

[14] Dainton (2007) p. 216

[15] Such views will be discussed in the next section of this thesis.

perceptual content under conjunction, and so in giving an analysis of closure, you get phenomenal unity for free.

In this sense, you might think of the subsumption view held by Bayne (and Bayne & Chalmers) as being a reductive view of phenomenal unity.

However, despite this, the contrast I am intending to invoke here is the contrast between those views which explain phenomenal unity in terms of some other kind of unity, and those views, such as the above, which do not. This is the reductive/non-reductive distinction I am drawing attention to, so I will continue to treat the subsumption view as a non-reductive view of phenomenal unity.[16] As it turns out, I will come to endorse an explanation of phenomenal unity that is not reductive in the sense of reducing phenomenal unity to some other kind of unity, but which is reductive in the sense that Bayne's view is.

Before moving on to discuss transitivity and other logical properties attributed to the phenomenal unity relation, I will here briefly return to the issue of the individuation of experiences at a time. This issue was raised in the previous chapter with respect to Tye's view of phenomenal unity, and is important to Bayne's view also. I will also look briefly at some objections to Bayne's mereological view and how he responds, pointing out that many of these objections are directed at not just Bayne's view, but all received views of phenomenal unity, and thus any view which treats phenomenal unity as a relation, primitive or not, that holds between individual perceptual experiences.

3.4 Individuating Experiences Revisited

Bayne's mereological view treats phenomenal unity as a relation that holds between token experiences. This it has in common with Bayne & Chalmers' joint view, and Dainton's view (despite the top-down vs., bottom-up difference). On Bayne's view, token experiences are phenomenally unified because they are parts of a larger experience, and states or experiences which are not parts of the same subsuming experience, such as my current auditory experience of the trumpet and *your* current auditory experience of it, cannot be phenomenally unified. This, points out Bayne, is at odds with some conceptions of how we should individuate experiences, notably the one-experience view of Tye's.

[16] Thanks to Tim Bayne for helpful discussion on this point.

I have already addressed this issue of the individuation of experiences in an earlier chapter of this thesis (see Chapter Two), but I will only bring it up again here insofar as it is relevant to our current concerns. The reason why raising this issue once again here is merited is because Bayne claims that the account of individuating experiences that he prefers ties in very naturally with his mereological account of phenomenal unity, and so it is worth considering if they perhaps stand or fall together, or whether an explanation of one of these issues has priority over the other.

The view of individuating experiences that Bayne advocates is called the *tripartite view*. According to this view, experiences are to be individuated according to subject of experience, time, and phenomenal properties. So, in order for experiences to differ, they must do so in terms of the subject of the experience, the time at which they are had, and the phenomenal properties instantiated by them. Here Bayne also points out that the account can also be thought of in terms of its concerning phenomenal *events* rather than properties, if we think of events in terms of the instantiation of properties.[17]

There is a "natural fit", claims Bayne, between this tripartite account of experiences, and his mereological conception of phenomenal unity. This natural fit comes from the fact that within a subject's overall phenomenal state at a time, there will be more or less complex experiences which we can identify. Bayne's example is of tasting a strawberry. The experience produced by tasting the strawberry will involve the instantiation of a range of phenomenal properties, such as 'tanginess' or 'sweetness'. According to Bayne's tripartite conception of individuation, we can think of these phenomenal properties as each involving distinct experiences which form part of the overall experience of tasting the strawberry, which in turn will form a part of the subject's overall phenomenal state at that time. Some of our more complex experiences will, Bayne claims, be modality-specific, such as your overall visual experience, and others will have content drawn from various sense modalities, such as your overall phenomenal state at a time, which will count as one complex experience according to Bayne's view.

The point to consider here is the following claim that Bayne makes: different approaches to the individuation of experiences might be appropriate in different contexts. This is coupled with the claim that there is a natural fit between the tripartite conception of experiences, and the mereological conception of phenomenal unity.[18] One thing to consider here is how helpful the tripartite conception of individuation really is

[17] I will, when discussing Bayne, use 'event' as he does. Thus a subject's overall phenomenal state at a time may also be described as an event.
[18] Bayne (2010) p. 24

when it comes to *synchronic* phenomenal unity, since the time will always be constant, and due to that the subject will also be fixed. This means that we will be individuating experiences solely according to the instantiation of phenomenal properties.

As I have suggested in Chapter Two of this thesis, it may be correct to say, as Bayne does, that when it comes to the individuation of perceptual experiences, especially the synchronic individuation of perceptual experiences, there may not be one uniquely privileged correct method of individuation, and so it may be permissible to divide up perceptual experiences in the manner according to your theoretical need. Tye divides up experiences in a way that fits with his theoretical needs, as we have seen, and Bayne makes the same move, advocating (though without claiming that it is the correct method) a view of the individuation of experiences which fits well with his mereological account of phenomenal unity.

One question that might be raised here, is how much our accepting a view of phenomenal unity (either Bayne's or Tye's) depends also on our being amenable to the view of the individuation of experiences that comes with it? Much opposition to Tye's 'one experience' view will stem from the fact that it burdens us with an unintuitive view of how many experiences we have at the one time (and indeed over time). The same thing might be case with those who are not favourably disposed to the tripartite view of the individuation of experiences. If we thought that the tripartite view of the individuation of experiences gave us counter-intuitive results, then we might be less inclined to accept Bayne's mereological account. It may then be the case that there could be pressure on Bayne's accounts, either of phenomenal unity or of the individuation of experiences, as those who advocate a different account of either may seek to drop one or other.

There are various ways that matters could play out here: if it turns out that there is only one plausible view of what is responsible for phenomenal unity, and if this view has implications for how we should individuate experiences, then this would count as a reason to adopt that view of individuation. However, if there are, as Bayne and I have both suggested, numerous ways of individuating experiences, no one of which is uniquely privileged, then we can adopt our preferred one, taking as some motivation to do so its fitting well with our preferred (or the correct) view of phenomenal unity. Pressure can come in both directions here then: those who hold one method of individuating experiences may claim an obligation to drop an account of phenomenal unity that fits ill

with it or vice versa, where a good model of phenomenal unity may bring with it pressure to drop some ways of individuating experiences.

It is not just Bayne that takes a liberal stance on the individuation of experiences. Dainton also allows that there are numerous equally legitimate ways that we can divide up a subject's overall experiential state at a time. The same question may apply to Dainton then, does his view of phenomenal unity depend on our adopting a certain method of the individuation of experiences at a time? And, if this preferred method is not presented as *the* correct method of individuating experiences, then what compels us to adopt the proposed view of phenomenal unity?

This issue will re-emerge again in the following chapter, with respect to Bayne's response to split-brain cases.

3.5 Objections to the Mereological Account

Here I will address some of the objections to Bayne's mereological account, and his replies, pointing out that the non-reductive views of phenomenal unity face the same objections we that we saw Tye level at the received views. Recall that received views, as Tye labels them, all involve some kind of relation that has to work to phenomenally unify separate perceptual experiences. I have addressed Tye's objections to the received views in the previous chapter, so I will go over them again only briefly here, along with Bayne's replies.

As we have seen, the first objection of Tye's, is a regress objection, which leads Tye to claim that the idea of a maximal, unifying state or experience is incoherent, and that Bayne's account is self-undermining.[19] If a subject has five phenomenally unified sense-specific experiences at a time, and what is responsible for their being unified is that they are all subsumed by a maximal experience, then there must be some sixth state which is itself an experience. This gives the five sense-specific experiences and the experience of unity which subsumes them. However, Tye claims, this maximal subsuming experience has to be an experience in its own right, as this will be the only way to explain why what it is like to have the maximal experience involves something over and above the character of the unified experiences. Bayne then describes Tye's first

[19] This objection of Tye's, and the others that Bayne considers here, were aimed originally at Bayne & Chalmers joint view (as an example of the received view of phenomenal unity), but since Bayne's own view is not significantly different, they will have purchase here too.

objection as follows (the following quotation is couched in Bayne's terminology rather than Tye's):

> Consider a maximal phenomenal state (e_m); that is, a phenomenal state that is not subsumed by any other phenomenal state. e_m includes as two of its elements a visual experience (e_1) and an auditory experience (e_2). Not only are these two experiences unified with each other, but each is unified with e_m. Now, if the unity of each of these two experiences requires that there be a unifying experience that subsumes them, then the unity of e_1 (e_2) and e_m seems to demand that the subject have a further experience, 'bigger' than e_m, that subsumes both e_m, and e_1 (e_2). But we stipulated that e_m was a maximal phenomenal state, a state not subsumed by any other phenomenal state. So if phenomenal unity is a relation between experiences, then the notion of a maximal phenomenal state is incoherent. But the notion of a maximal phenomenal state clearly *is* coherent, so phenomenal unity cannot be a relation between experiences.[20]

Bayne's reply to this objection is to block the move which leads to the maximal state becoming incoherent. The move in Tye's argument that Bayne rejects (and which any received view of phenomenal unity which posits a maximal unifying state would reject) is that which requires the state which unifies e_1 and e_m to be 'bigger' than e_m itself. Bayne's claim is that nobody who endorses the mereological view need endorse this. e_1 and e_2 are unified with e_m, but given that parthood is a reflexive relation, e_m is also unified with itself. This allows the proponent of the mereological view to hold that what unifies the maximal phenomenal state with its parts is nothing other than the maximal state itself. Hence, there is no reason why the mereological account has to deny that the idea of a maximal phenomenal state is incoherent, nor do we embark on a vicious regress of unifying states.

As I have outlined before, Tye has another regress-based objection which is directed towards the mereological view. Bayne here calls this objection the *phenomenal bloat* objection. It has as its target the idea that the phenomenal unity relation must itself be the object of experience. If the phenomenal unity relation did not possess its own phenomenology, claims Tye, then there would be nothing it is like to have phenomenally unified perceptual experiences. Thus, since the phenomenal unity relation must have its own phenomenology, this must be unified with the subject's experiences, and now this necessitates a further phenomenal unity relation. This concludes Tye, leads to a vicious

[20] Bayne (2010) p. 29

infinite regress. Bayne formulates this argument as follows (again using his own terminology rather than Tye's):

(1) Suppose that phenomenal unity is a relation between experiences $e_1 – e_5$ (assumption for *reductio ad absurdum*).

(2) This unity relation (R_1) between experiences must itself be experienced, for if there were no experience of the unifying relation, then there would be nothing it is like for $e_1 – e_5$ to be unified. (Alternatively, there would be no phenomenal difference between a situation in which $e_1 – e_5$ were phenomenally unified and a situation in which they are not unified.)

(3) If R_1 is itself experienced, it must have its own phenomenology.

(4) If R_1 has its own phenomenology, its phenomenology must be unified with that of $e_1 – e_5$.

(5) In order to account for the fact that R_1 is unified with $e_1 – e_5$ we need to posit another unity relation (R_2).

(6) But of course R_2 must itself be experienced, for if there were no experience of the unifying relation, then there would be nothing it is like for $e_1 – e_5$ to be unified with R_1.

(7) But now we have embarked on a vicious infinite regress.

(8) So (1) must be false: phenomenal unity is not a relation between experiences.[21]

After noting that this argument threatens to take out not just his own mereological account of phenomenal unity (and other received view accounts), but all accounts of phenomenal unity. Bayne gives his favoured response to this objection. I have mentioned this response earlier in this thesis, but will explain it once more here in brief.

Bayne's response to this argument of Tye's is to reject the premise which states that in order to make a phenomenal difference, the phenomenal unity relation must itself have a particular phenomenology which is experienced. As I have mentioned previously, Bayne's response is to say that phenomenal unity is a phenomenal relation insofar as it

[21] Bayne (2010) p. 30

makes a difference to a subject's overall phenomenology, but not in the sense that it is itself an object of experience.[22]

Two more objections to Bayne's view that he outlines and which I will consider briefly come from arguments by Andrew Brook and Paul Raymont, and John Searle respectively.[23] The Brook and Raymont objection is in the spirit of a point made by William James. The crux of the claim is that the mere combination of experiences does not automatically give us the experience of a combination. Bayne's response to this objection is to say that we should not expect an account of phenomenal unity to give us this kind of integration. The mereological account, claims Bayne, is only an account of how it is that our experiences are phenomenally unified, that is, how they get to have conjoint phenomenology. It is not an account of how it is that a subject's experiences become representationally integrated. For example, it is not an account of how it is that the parts of a face are experienced as a coherent whole. Nor is it an account of how the different features or properties of an object are experienced as all inhering in that object. A theory of phenomenal unity does not necessarily give us a theory of representational unity, and as Bayne notes here, though representational unity and phenomenal unity are closely connected, they should not be identified.[24] One interesting upshot of making this clarification is that putative cases of disunity in consciousness within a subject may be breakdowns in some form of representational unity, rather than a breakdown in phenomenal unity. I will return to this point when further discussing the possibility of breakdowns in phenomenal unity.

The final objection to the mereological view of phenomenal unity that Bayne considers comes from John Searle, who objects to mereological conceptions of consciousness full stop.

> The urge to think of consciousness… as made up of smaller building blocks is overwhelming. But I think it may be wrong for consciousness… Indeed, maybe it is wrong to think of consciousness as made up of parts at all… Instead of thinking of my current state of consciousness as made up of the various bits – the perception of the computer screen, the sound of the brook outside, the shadows cast by the evening sun falling on the wall – we should think of all of these as modifications, forms that the

[22] The same kind of response is made by Barry Dainton (2006) & (2007) .
[23] Brook & Raymont (2009) and Searle (2000).
[24] There are of course potential reductive explanations of phenomenal unity which do identify phenomenal unity with some kind of representational unity. I will discuss such a view in Chapter Five.

> underlying basal conscious field takes after my peripheral nerve endings have been assaulted by the various external stimuli.[25]

In response to this scepticism about 'building block' approaches to consciousness, Bayne takes care to suggest that his mereological account of phenomenal unity is not the kind of account the Searle is objecting to. The building block view suggests that the parts of the overall conscious field are in some sense prior to the whole, and perhaps even the stronger claim that the parts are the 'real' units of consciousness. This seems to be what Searle is objecting to, but Bayne is at pains to point out that his mereological view takes no such stance on the explanatory (or otherwise) priority of the parts over the whole. In fact, given that Bayne advocates a top-down view of phenomenal unity which may grant explanatory priority to a subject's overall experience at a time, then this further supports the claim that his mereological view should not be the target of Searle's objections.

We have seen in this short section that Bayne considers the objections to received view positions on phenomenal unity raised by Tye and other, and has replies to them which may be adopted by any other received view theorist, no matter if they explain phenomenal unity reductively or non-reductively. As I have mentioned previously, these objections apply to Bayne's mereological view, as despite its positing of a primitive relation (along with Bayne & Chalmers and Dainton), it still shares with other received view positions the central tenet that phenomenal unity is a relation that serves to unify several perceptual experiences, and as such is as opposed to Tye's view of phenomenal unity as much as any other received view.

To further clarify things at this stage: the *received view* is a view of what phenomenal unity *is*. This view states that phenomenal unity is some kind of relation which holds between token perceptual experiences and unifies them. The received view however is neutral on how phenomenal unity is explained further.

Non-reductive views are views about how to *explain* phenomenal unity. They hold that phenomenal unity should not be explained in terms of some other kind of unity that consciousness exhibits. Some non-reductive views are also received views, such as Dainton's, Bayne's and Bayne & Chalmers'. However, other non-reductive views are not also received views, such as Tye's. Since I have argued that we should adopt a received view of what phenomenal unity is, this chapter is dealing with received, non-reductive views of how to explain phenomenal unity.

---

[25] Searle (2000) p. 575

I will move on now to discuss some of the logical features attributed to the phenomenal unity relation by non-reductive views. The most important of which will be transitivity.

3.6 Reflexivity, Symmetry and Transitivity

Phenomenal unity is basic, *primitive relation*. This is the central claim that Dainton makes. Phenomenal unity is a subsumption relation. This is the view of Bayne and Bayne & Chalmers. We may also think of this subsumption relation as some kind of primitive. In order to begin to unpack these claims further, we need to know what *kind of* relation phenomenal unity is, and what more we can say about it, further to the claims above. I have in the previous sections given some detail on how Bayne & Chalmers and Bayne see the subsumption relation, and how Bayne's mereological account works. I will now move on to discuss the other properties that are attributed to the phenomenal unity relation by non-reductive views.

Phenomenal unity is on Dainton's view a material relation, as it is a relation between experiences, which he treats as concrete particulars. Second, because phenomenal unity is a relation which holds between distinct experiences, we might initially think that it has to be a dyadic or two-place relation. However, as Dainton imposes no upper limit on the number of experiences which can be phenomenally unified, it will not always be a relation of the same degree. Further, given that we can divide experiences into parts, and all those parts will always, or at least usually, be phenomenally unified with each other, it makes better sense, according to Dainton, to say that experiences are phenomenally unified with themselves. This makes phenomenal unity a *reflexive* relation. We can take the other non-reductive views of Bayne and Bayne & Chalmers to hold this claim about reflexivity also, as they invoke it in replying to Tye's regress objections, as we have seen. Phenomenal unity's being reflexive allows a subject's overall experience at a time to be self-unifying, with no need for another state which unifies it with its contents.

The second key property that Dainton attributes to phenomenal unity is that it is *symmetrical*. If experience e_1 is unified with experience e_2, then e_2 will automatically be unified with e_1. Due to their analysis of phenomenal unity in terms of subsumption however, Bayne & Chalmers, and Bayne, do not hold this claim also. Though the maximal state e_m may subsume e_1, the converse does not hold. Subsumption *is not* a

symmetrical relation. This is a point of departure between Dainton's primitive relation view on the one hand, and the subsumption view of Bayne and Bayne & Chalmers on the other. However, despite this difference, little will turn on this point.

Dainton's final claim is that phenomenal unity is *transitive*. So, if a subject has three experiences, e_1, e_2 and e_3, and if e_1 is phenomenally unified with e_2 and e_2 with e_3, then by the law of transitivity, e_1 and e_3 will also be unified.

Since the phenomenal unity relation, according to Dainton's position, is reflexive and symmetrical, if it is transitive also, it will be an equivalence relation. Equivalence relations partition their relata into discrete groups, which do not overlap. Further, it should be mentioned that equivalence groups exhaust the domain. If phenomenal unity is an equivalence relation, then the totality of a subject's experiences at any given time will be divided into discrete, non-overlapping streams of consciousness (again, I am only referring to the synchronic here). It might be asked here 'if we have two discrete, non-overlapping streams of consciousness, in what sense do we have a single subject?'. Here, much depends on the conditions imposed on a subject. as I mentioned in the first chapter, there are various methods by which we might individuate subjects, and as I have also said above, it is not my intention to defend any view here. I will simply not that there are views on which a single subject can have two-separate streams of consciousness.[26]

This view of the phenomenal unity relation as transitive would if true give a very neat and clean picture of streams of consciousness, and what they would look like at a time, and further would give us a potential method by which we could determine whether or not states are part of the same stream of consciousness. Thus if phenomenal unity is an equivalence relation, it would be of epistemological benefit also. If the phenomenal unity relation were to be an equivalence relation, it would also gives us a picture of how the phenomenal unity relation, and streams of consciousness, are related, and what the role of phenomenal unity is: it is the relation which, if it holds between states, ensures that those states are part of the same stream of consciousness, and means that they will be phenomenally unified with all the other states in that stream of consciousness. Further, as we shall see, if the phenomenal unity relation is *transitive*, then this will mean that phenomenal unity cannot break down within a single subject. This will be a key point which I will explain shortly.

[26] See for example Dainton's own (2006), where he considers split-brain patients to be a single subject with two streams of consciousness.

However, since Bayne and Bayne & Chalmers do not treat phenomenal unity as symmetrical, they do not hold that it is an equivalence relation. A further blow to this view will be dealt in the next chapter, where I will try to cast doubt on any transitivity claims by looking at potential counterexamples.

Though they deny that it is symmetrical, Bayne, and Bayne & Chalmers, also hold that synchronic phenomenal unity is transitive. Though all parties suggest that phenomenal unity may not be transitive with respect to phenomenal unity *over time*. As an indication of his commitment to the transitivity claim, Bayne says the following (after noting that *diachronic* transitivity failures are possible):

> [I]t is far less plausible to suppose that transitivity can fail for sets of *simultaneous* experiences. In fact, it is tempting to suppose that the phenomenal field *cannot* fragment in the way that a failure of transitivity would require. In other words, it is tempting to suppose that for any three simultaneous experiences, e1,e2 and e3, if both1 and e2 are phenomenally unified with e3 then they must also be unified with each other. Let us call the assumption that phenomenal unity is transitive with respect to simultaneous states *the transitivity thesis.*[27]

If it turns out to be the case that the phenomenal unity relation is not transitive, then there is the possibility of partially phenomenally unified streams of consciousness. Partially phenomenally unified streams of consciousness would involve states in a stream of consciousness that were not phenomenally unified with the rest of the states in that stream.

Examples of this kind of partial phenomenal unity might be the kind of faint bodily sensations that linger at the periphery of one's consciousness, such as the slight back pain I feel whilst sitting at my computer in an uncomfortable chair: at a time, this sensation, though there is something it is like for me to undergo it, may nevertheless not be experienced together with the rest of my perceptual experiences.[28] This case would give us a stream of consciousness belonging to a subject where all bar one of the subject's states were phenomenally unified with each other, and that one state would not be unified with any other state. A more plausible example of this kind of partially unified stream of consciousness would be one in which the back pain experience was phenomenally unified with the subject's other bodily experiences, but not with their

[27] Bayne (2010) p. 37
[28] Dainton (2006) p. 90

auditory experiences, despite the auditory experiences and the other bodily experiences being phenomenally unified. These cases however, seem hard to imagine from a first-personal standpoint. A more striking example of partial phenomenal disunity, and an initial candidate for this kind of case, is a split-brain subject. These cases will be dealt with in the next chapter, but for now I will outline some more features of the non-reductive views I have been looking at in this chapter.

Holding this transitivity thesis (on which phenomenal unity is transitive with respect to simultaneous experiences) is part of the motivation for another thesis that Bayne and Bayne & Chalmers hold, *the unity thesis*.

3.7 The Unity Thesis

The Unity Thesis is outlined in Bayne & Chalmers (2003) paper, but here I will stick with the formulation in Bayne (2010). The central tenet of the thesis is the same in both cases. The unity thesis is formulated as follows:

> Unity Thesis: Necessarily, for any conscious subject of experience (S) and any time (t), the simultaneous conscious states that S has at t will be subsumed by a single conscious state – the subject's total conscious state.[29]

The kind of unity that this thesis refers to is phenomenal unity. Though this unity thesis does not fall out of Bayne's account of phenomenal unity, it does depend for its truth on the truth of the transitivity thesis, which Bayne holds, and so stands or falls on the truth of an explanation of phenomenal unity which makes the same commitments to transitivity.

To spell out the relation between the unity thesis and the transitivity thesis in more detail: the transitivity thesis, as we have seen, entails that synchronic phenomenal unity is a transitive relation. The unity thesis, which if true means that any two states of a single subject are necessarily phenomenally unified, depends on the transitivity thesis being true, as for it to be the case that any two states of a subject are necessarily unified, there cannot be any failures of transitivity in any cases. For a subject with three perceptual states for example, if the phenomenal unity relation were not transitive, then it would be

[29] Bayne (2010) p. 16

possible for two of these three states to not be phenomenally unified, and this would falsify the unity thesis.

It is the necessity claim of the unity thesis which requires transitivity of phenomenal unity. If the unity thesis stated only that in normal circumstances, two states of a single subject will be phenomenally unified, then this could be true even if the transitivity thesis turned out to be false.

Two further points should be made here about the relation between the unity thesis and the transitivity thesis: firstly, the unity thesis needs various other things to secure its truth, aside from transitivity. For example, the defender of the unity thesis needs the 'two-streams' model of split-brain subjects to be false, as we will see in much greater detail in a later chapter. Second, though the unity thesis requires the truth of the transitivity thesis, the transitivity thesis does not itself entail the unity thesis. This also will be discussed a propos of split-brain cases in a later chapter.

One feature of the unity thesis which merits further discussion at this point is what kind of necessity is being invoked here. Bayne (and Bayne & Chalmers) will argue that we never have phenomenally disunified consciousness. However, as Bayne explains it, the unity thesis is not as strong as it might seem. Bayne claims that the unity thesis should not be taken as putting forward a conceptual or metaphysical truth about phenomenal unity: it is not claiming that phenomenal unity holds in all metaphysically possible worlds, nor is it claiming that it is part of the concept of phenomenal unity that it always holds. Indeed Bayne goes further than this, and suggests that it is not even nomological necessity that is being invoked in the unity thesis

> I do not even claim that the unity of consciousness is grounded in the laws of nature. Perhaps there are surgical innovations or evolutionary developments that could bring about a division in the stream of consciousness; perhaps there are other species in which the unity of consciousness can be lost. My only claim is that we have no good reason to think that any such division has actually occurred in the members of our own species.[30]

The unity thesis then, states only that there has never been a case of phenomenal disunity in a single human subject. We might ask of the unity thesis then, in what sense it involves a necessity claim at all.

I will return to the unity thesis and to the transitivity thesis in the next chapter. Before that I will examine where non-reductive explanations of phenomenal unity stand

[30] Bayne (2010) p. 17

on the issue of whether or not phenomenal unity involves extra phenomenology, over and above the phenomenology of the unified states. That is, as an explanation of what is responsible for phenomenal unity, do non-reductive accounts have any commitments to what phenomenal unity *is*?

3.8 Does 'Conjoint' mean 'Extra'?

We have seen from the discussion in Chapter One, that phenomenal unity is standardly taken to involve conjoint phenomenology. This conjoint phenomenology is what it is like to have two or more states simultaneously and for them to be phenomenally unified. We have also seen that it is not clear whether or not conjoint phenomenology means extra phenomenology.

Of the views I have outlined, do any make any claims on this issue? Here is Bayne's description of phenomenal unity, which I introduced in the first chapter

> Consider again what it's like to hear a rumba playing on the stereo whilst seeing a bartender mix a mojito. These two experiences might be subject unified insofar as they are both yours. They might also be representationally unified, for one might hear the rumba as coming from behind the bartender. But over and above these unities is a deeper and more primitive unity: the fact that these two experiences possess a *conjoint experiential character*. There is something it is like to hear the rumba, there is something it is like to see the bartender work, and there is something it is like to hear the rumba *while* seeing the bartender work. Any description of one's overall state of consciousness that omitted the fact that these experiences are had together as components, parts, or elements of a single conscious state would be incomplete. Let us call this kind of unity – sometimes dubbed 'co-consciousness' – *phenomenal unity*.[31]

And again, Bayne & Chalmers' description:

> [T]wo states are phenomenally unified when they have a conjoint phenomenology: a phenomenology of having both states at once that subsumes the phenomenology of the individual states. When A and B are phenomenally conscious states, there is something it is like for a subject to have A, and there is something it is like for a subject to have B. When A and B are phenomenally unified, there is not just something it is like to have

[31] Bayne (2010) p. 11

> each state individually: there is something it is like to have A and B together. And the phenomenology of being in A and B together will carry with it the phenomenology of being in A and the phenomenology of being in B.[32]

As I said when introducing these views in Chapter One, there seems nothing in these descriptions which would commit either view to the claim that conjoint phenomenology involves extra phenomenology, but also nothing that would prevent them from holding that position. Recall also that in replying to Tye's regress objections, both Dainton and Bayne, proponents of non-reductive views, hold that the phenomenal unity relation, despite not being itself an object of introspection, does make a phenomenal difference. This kind of claim may suggest that the phenomenal unity relation is responsible for some extra phenomenal character that is not attributable to any one of the unified states, though again, it is not conclusive.

In either case, if we were to suppose that conjoint phenomenology involves extra phenomenology; would a non-reductive explanation of phenomenal unity allow us to readily explain this extra phenomenal character?

Again here, much will depend on how you see the relation between content and character. If you hold that for there to be extra character, there must be extra content, then we will need an account of how the phenomenal unity relation as described by the non-reductive views can provide extra content, and what kind of content it would be. On the face of it, such an account seems as though it would be difficult to provide. This may suggest that non-reductive views of phenomenal unity which posit extra phenomenal character would not fit will with representationalist views of the relation between character and content.

Non-representationalist, non-reductive views may be able to more plausibly posit extra phenomenal character, as they can posit changes in character without any change in content. Perhaps it might be argued that it is just a basic feature of experience that the phenomenal unity relation brings with it some extra phenomenal character that is over and above the sum of the character of the individual unified states.

As an aside here, there is an interesting question concerning whether any extra character instantiated by the phenomenal unity relation must be *amodal* in nature. Amodal in this sense means not associated with any one sensory modality in particular. We have seen that phenomenal unity can be described in the following way:

---

[32] Bayne & Chalmers (2003) p. 32

> At a time, a subject's overall phenomenal state is such that it involves phenomenal character which is over and above the sum of the character of the individual phenomenal states the subject has at that time.

This conjoint phenomenology is instantiated when the subject's phenomenal states are phenomenally unified. Phenomenal unity may necessarily involve the subject's overall phenomenal state possessing extra content (and extra character). This extra content is not associated with any particular sensory modality. And thus, any extra character is not associated with any particular experience the subject has at that time. Because of this, we might think that this extra content must be amodal content. And thus the extra character must be amodal character. Even if phenomenal unity involves extra character without extra content, the extra character may not be associated with any one of the unified experiences. I do not have the space however, to pursue this question further, and since I am not suggesting that any of the views I have looked at so far do in fact hold that conjoint phenomenology necessarily involves extra phenomenology, I will leave this question here. I will however, raise the issue of conjoint and extra phenomenology again, with respect to the reductive explanations of phenomenal unity that I will consider in the next section of this thesis.

3.9 Conclusion

To recap what has been outlined in this chapter then, I have introduced the idea that there are some views of phenomenal unity on which it is not to be explained in terms of some other feature of consciousness, or some other kind of unity that consciousness exhibits. Such a view is held by Dainton, Bayne & Chalmers and Bayne. I have tried to give an outline of the key points of a non-reductive view of phenomenal unity, and of the commonalities and differences between all three of the above views. All three of these views are received views of phenomenal unity, holding as they do that a subject has several perceptual experiences at a time which are unified. The differences between these non-reductive views are that Dainton takes phenomenal unity to be a bottom-up relation, and as a primitive. Both Bayne and Bayne & Chalmers hold that phenomenal unity is a top-down relation, with the starting point being a subject's overall experience at a time, rather than individual experience. As I have said however, little turns on the distinction

here, as neither side seem to advance any kind of priority claims here. Bayne and Bayne & Chalmers also give their analysis of phenomenal unity in terms of subsumption, rather than leaving the phenomenal unity relation as a bare primitive. I have discussed above however, that given all of these views hold that phenomenal unity should *not* be explained by reducing it to another form of unity, I will continue to hold that they are all non-reductive views in this sense.

The crucial point in common between all three of these views (though as I have pointed out, this need not be a feature of all non-reductive views) is their commitment to *the transitivity thesis*, which states that with respect to synchronic phenomenal unity, the unity relation is transitive. We also have seen how this motivates Bayne & Chalmers and Bayne's *unity thesis*, which states that for any two or more experiences had by a single subject at a time; these experiences will be subsumed by a single maximal state. Holding the unity thesis and the transitivity thesis essentially commits you to holding that phenomenal unity can never break down within a single subject. Though these commitments are not intrinsic to non-reductive views, they are made by all of the non-reductive views that I am considering, and so I will continue to investigate these commitments.[33]

[33] The only extant non-reductive (in the sense I am using here) view of phenomenal unity which does not make the transitivity claim comes from Tye (2003). I have already discussed Tye's view at length however, and so will not do so further, except to mention his view of split-brain cases in the next chapter.

Chapter Four – Non-Reductive Explanations and Split-Brain Cases

4.1 Introduction

This chapter will essentially be an assessment of how the non-reductive views of phenomenal unity outlined in the previous chapter account for split-brain cases, which seem initially to falsify the transitivity thesis by posing counterexamples. I will also at the close of this chapter offer some supplementary reasons for not adopting a non-reductive account of synchronic phenomenal unity.

After re-capping the transitivity and unity theses, I will detail the evidence which serves as the basis for claims that the split-brain cases involve some kind of breakdown in phenomenal unity. There are two kinds of account on which split-brain cases involve a phenomenal unity breakdown; these are *two-streams accounts* and *partial unity accounts*. It should be noted here that the two-streams approach, though it may falsify the unity thesis, does not necessarily threaten the transitivity thesis. I will then explain precisely the problems split-brain cases pose for Dainton, Bayne & Chalmers and Bayne.

After this, I will look at how the non-reductive views outlined in the previous chapter respond to the split-brain cases. Dainton adopts a two-streams approach, meaning he must drop the unity thesis, though as I will point out, the two-streams approach allows him to maintain the transitivity thesis. Bayne & Chalmers attempt to preserve the unity thesis also, and I will outline their joint attempt to deal with the split-brain cases, before looking at Bayne's more recent "switch model" of the split-brain cases. At the close of this chapter, I will also suggest that adopting an alternative account of phenomenal unity would free us from having to adopt a model which captures what is going on in all split-brain cases. Further, I will suggest that there is a general problem with explanations which rely on the positing of primitive relations, as their explanatory power is less than satisfying. We should look instead for accounts of phenomenal unity which can give a fuller, more reductive explanation (though I don't here exclude the views of Bayne and Bayne & Chalmers from this bracket).

4.2 Transitivity & the Unity Thesis

To begin this chapter, I will re-introduce the transitivity and the unity theses held by Dainton, Bayne & Chalmers and by Bayne himself. The transitivity thesis is put as follows:

> [I]n fact, it is tempting to suppose that the phenomenal field *cannot* fragment in the way that a failure of transitivity would require. In other words, it is tempting to suppose that for any three simultaneous experiences, e_1, e_2 and e_3, if both e_1 and e_2 are phenomenally unified with e_3 then they must also be unified with each other. Let us call the assumption that phenomenal unity is transitive with respect to simultaneous states *the transitivity thesis.*[1]

And the unity thesis

> Unity Thesis: Necessarily, for any conscious subject of experience (S) and any time (t), the simultaneous conscious states that S has at t will be subsumed by a single conscious state – the subject's total conscious state.[2]

As I have mentioned previously, it seems that the unity thesis, which is held by Bayne & Chalmers and Bayne gets support from the transitivity thesis. If there are no transitivity failures of phenomenal unity, then all of a subject's perceptual states at a time can always be subsumed by a single state. As I have also mentioned in the previous chapter, this unity thesis applies to phenomenal unity only. All parties here agree that other kinds of unity *can* break down, and may do so in split-brain cases.

4.3 Phenomenal Unity Breakdowns and Partial Unity

Before detailing the various responses that can be made to the split-brain cases, I will here present the case for thinking that these cases involve some breakdown in phenomenal unity. I will do this by presenting the partial unity and two streams models of the split-brain cases. As mentioned above, both of these models attribute some breakdown in phenomenal unity to split-brain subjects, at least under experimental conditions, though they present different pictures of how this disunity manifests itself.

[1] Bayne (2010) p. 37
[2] Ibid. p. 16

The proponents of the transitivity and unity theses need both to diffuse the disunity accounts, and provide a positive account of what is going on in split-brain cases which does not involve phenomenal disunity of any kind.

Split-brain cases, as we have seen in Chapter One, involve subjects who have had their corpus callosum severed. These subjects exhibit various dissociations under experimental conditions, which may be seen as evidence for some kind of disunity in consciousness, potentially phenomenal disunity. Take a standard split-brain test, as described in Chapter One: the subject, S, is presented with the composite word 'key-ring', so that 'key' falls within the subject's left visual field, and 'ring' falls within their right visual field. Since the visual system has a contralateral structure, the information presented to the left visual field will be processed by the right hemisphere, and information presented to the right visual field will be processed by the left hemisphere. S will report having seen only the stimulus which has been displayed on the right half of the screen ('ring') and further, will deny having seen the stimulus presented to the left visual field. At the same time, S will use his left hand to correctly retrieve the object matching the word presented to the left visual field ('key'). When asked to name the object selected by their left hand, S will respond "ring", matching the stimulus to the right visual field.

The seeming upshot of this kind of split-brain experiment is that S cannot jointly report 'key-ring', leading to the view that the reason for this is some form of breakdown in the unity of S's consciousness. One possibility then is that split-brain cases involve a breakdown in phenomenal unity. In the case above, S has an experience of 'key' and an experience of 'ring', but no experience of 'key-ring'. So, S has two experiences which do not possess conjoint phenomenology, and thus S's two experiences (as of 'key' and 'ring') are not phenomenally unified.

Before going into the particulars of the two-streams and partial unity models of phenomenal disunity, I will point out a common feature to these disunity models, highlighted by Tim Bayne.[3] Bayne claims that all arguments against the necessity claim he makes take the form of an argument by counterexample. The argument has two components: first a claim that the subject in the counterexample has two or more states which are conscious, and second, a claim that these states are not phenomenally unified. The second part is the important one for our purposes, and Bayne calls this the "negative moment" of the argument. He takes there to be two ways that the negative moment can

[3] Bayne (2010) pp. 194 - 199

plausibly be developed, by inference *from representational disunity to phenomenal disunity*, or *from access disunity to phenomenal disunity*. I will briefly examine both of these strategies.[4]

Representational disunity is one potential way that we could get evidence of phenomenal disunity. The kind of failures of representational integration that we would need could be found in the split-brain cases. For example, suppose that e_1 and e_2 are representations of the first and second halves of the word 'cobweb', i.e. e_1 is 'cob' and e_2 'web'. If e_1 and e_2 are phenomenally unified for a subject, and if that subject recognises the word 'cobweb', then they will typically enjoy an experience as of the word 'cobweb'. If $e_1$and e_2 are not phenomenally unified, then the subject will have experiences of the words 'cob' and 'web', without experiencing it as 'cobweb'. It is from these kinds of cases that we may be able to extract evidence for phenomenal disunity, suggests Bayne. He further suggests that these arguments will require some principle connecting representational and phenomenal unity, which he outlines as follows:

> *Representational Integration Principle* (RIP): For any pair of simultaneous experiences e_1 and e_2, if e_1 and e_2 are phenomenally unified then, *ceteris paribus*, their contents will be available for representational integration.[5]

For this principle to do any work in providing evidence of phenomenal disunity, we are required to show that the ceteris paribus clauses cannot be activated. Bayne issues several cautionary points here: firstly, some creatures and subjects will not, for whatever reason, have the ability to integrate the contents of their experiences in the appropriate matter. We should thus be wary of tying phenomenal unity to the possession of certain integrative capacities. Second, even if the subject in question possesses the relevant capacities, they might be prevented from exercising them in the instance in question, so we should be cautious when arguing from representational disunity to phenomenal disunity using a case which does not involve a subject in normal attentive wakefulness.

In light of these cautionary notes, Bayne urges us to be careful, and not to suppose that representational unity and phenomenal unity are tied together with strict necessity.[6]

Another way to potentially provide evidence of phenomenal disunity is to argue from *access disunity*. This kind of disunity may also be found in split-brain cases. The subject in a case like this would be undergoing two experiences at a time, e_1 and e_2, but

[4] Bayne (2010) pp. 197-199
[5] Ibid. p. 106
[6] Ibid. p. 107

without e_1 and e_2 being available to the same consuming systems. e_1 might be available for verbal report but not for reasoning, and e_2 might be available for use in reasoning, but not for verbal report. This breakdown in access unity may be taken as evidence of phenomenal disunity. Again, a linking principle would be needed here, between access and phenomenal unity, and Bayne puts this principle as follows:

> *Conjoint Accessibility Principle* (CAP): For any pair of simultaneous experiences e_1 and e_2, if e_1 and e_2 are phenomenally unified then, *ceteris paribus*, their contents will be available to the same consuming systems.[7]

What is it for two experiences to be access disunified? Bayne describes it as involving a subject's experiential states failing to be available to the same consuming systems of that subject.

> Take a creature with two experiential states (e_1 and e_2) and five consuming systems (CS_1 … CS_5). States e_1 and e_2 will be fully access unified if their contents are available to all and only the same consuming systems, and fully access *dis*unified if their contents are not available to any of the same consuming systems. But suppose that the contents of e_1 and e_2 are available to *some* of the same consuming systems but not others… We might think of such states as *partially access unified*.[8]

If it is possible for the contents of a subject's states to be available to some consuming systems but not others, then access unity and disunity can come in degree. This is something which may be resisted, for various reasons, amongst which would be the desire to preserve the unity thesis. Bayne however, does not think that partial access disunity can be dismissed easily.[9]

Given this, Bayne claims that the most sensible thing to say here is that "the strength of any argument from access disunity will be a function of the degree to which the contents of the relevant states are co-accessible: the less co-accessible the contents of the states, the stronger our evidence for thinking that they are not phenomenally unified".[10]

---

[7] Bayne (2010) p. 108
[8] Ibid.
[9] Ibid. p. 109
[10] Ibid.

Bayne & Chalmers do not think that access disunity entails phenomenal disunity (though I will go on to question their argument as to why this is so), and Bayne takes this line also. As with the argument from representational disunity, Bayne thinks that there are various ways in which things will not be equal and so the ceteris paribus clauses in CAP will not get off the ground.

The first way in which things might not be equal is due to the presence of processing bottlenecks. Arguments from access disunity to phenomenal disunity will be most plausible when dealing with conscious states which are roughly the same 'size' in terms of their contents. The greater the size disparity, the more chance of the failures of access unity being due to some processing bottleneck rather than a breakdown in phenomenal unity. This is the line of argument used by Bayne & Chalmers. They claim that the seeming access disunity in the Sperling experiment is due to a processing bottleneck rather than a breakdown in phenomenal unity, and we should think the same thing of split-brain cases.[11] I will examine this response to the split-brain cases in greater detail later in this chapter and will put pressure on this claim by pointing out the failure of the analogy between the Sperling experiment and the split-brain cases. This of course does not refute Bayne here, but it does suggest that as well as being cautious in our use of arguments from access disunity to phenomenal disunity, we should also be cautious of drawing analogies between all cases where processing bottlenecks are present.

It is not simply the presence of processing bottlenecks that may stop things being equal however. Things may not be equal because of differences in the representational format of the states in question. If, of two states of a subject, one has conceptual content and the other only non-conceptual content, then, Bayne points out, given a plausible link between conceptual content and reasoning, it would be no surprise if the state with conceptual content was able to interact with belief-revision mechanisms in a way that the state with purely conceptual content was not.[12]

> The lesson to be learnt from this is that failures of access unity will provide better evidence of failures in phenomenal unity when the states in question share a common representational format. Discovering that e_1 and e_2 couldn't (say) be jointly reported would give us a better reason to think they are not phenomenally unified if they were

[11] Bayne & Chalmers (2003)
[12] Bayne (2010). p. 110

> both conceptual states as opposed to one of them being conceptual and the other non-conceptual.[13]

So, despite there being reason to think that phenomenally unified states will generally be both representationally and access unified also, Bayne urges caution when we formulate arguments based on cases of access or representational disunity. When examining how Bayne's mereological model would deal with the split-brain cases I will discuss, I will bear this caution in mind.

A final point I will note here before looking at the split-brain cases once again, concerns the notion of *probe-dependence*. A great many things that we study are independent of our studying them: the number of people in a building, or the number of planets in a solar system, are independent of our measuring them. If consciousness is not independent of our studying it in this way, then it may be said to be 'probe-dependent'.[14] Consciousness being probe-dependent would mean that certain responses on behalf of a subject, which we would think of as being indicative of some feature or other of consciousness, may not be independent of the probe we have used to measure the response in question.

> The experiences e_1 and e_2 might appear to be available to different consuming systems, but that appearance might be an illusion generated by the probe-dependence of consciousness. Suppose that we show a subject a light at the same time as we play a sound to him. We want to know whether he experiences both the sound and the light, and – if so – whether or not these two experiences were phenomenally unified. How are we to test him? Requiring him to produce a verbal report might bias him to report the light (and perhaps also extinguish the sound), but requiring him to produce a button-press report might bias him to report the sound (and perhaps also extinguish the light). We might be tempted to think that these two experiences were simultaneous but not phenomenally unified with each other, but in fact there may have been no single trial on which the subject was simultaneously aware of both the light and the sound.[15]

[13] Bayne (2010) p. 110
[14] Ibid. p. 112
[15] Bayne (2010) p. 115

If consciousness is genuinely probe-dependent, and Bayne gives some evidence to suggest that it is, then arguments from access disunity will have to be extra vigilant to any bias caused by the probe employed in any given case.[16]

Despite the caution Bayne urges above, he does concede that strong arguments can be formed using the above tactics and inferring phenomenal disunity from the presence of access or representational disunity. To take the latter first, phenomenal disunity is inferred from the lack of representational integration between the left and the right hemispheres, using the following kind of argument which Bayne formulates thus:

> (1) S [a typical split-brain patient in a key-ring experiment] has, simultaneously, and experience with the content <'key'> and an experience with the content <'ring'>.
> (2) Any subject with simultaneous experiences of <'key'> and <'ring'> that are phenomenally unified with each other will also have an experience with the content <'key' & 'ring'>.
> (3) S does not have an experience with content <'key' & 'ring'>.
> (4) S's experiences of 'key' and 'ring' are not phenomenally unified.[17]

Bayne concedes that this argument seems secure. Though the truth of premise (2) seems to rely on the representation integration principle (RIP) that is discussed above, Bayne seems happy that the ceteris paribus clauses do not present any problems in this case.

Using the key-ring experiment, it is also possible to construct an argument from access disunity to phenomenal disunity. Bayne formulates this argument as follows:

> (1) S has, simultaneously, an experience with the content <'key'> and an experience with the content <'ring'>.
> (2) If simultaneous experiences of <'key'> and <'ring'> are phenomenally unified with each other then they will be access unified: their contents will be available to the same range of consuming systems.
> (3) S's representations of 'key' and 'ring' are not access unified: although the contents of both states are available for high-level consumption, they are not available to the *same* consuming systems.
> (4) So, S's experiences of 'key' and 'ring' are not phenomenally unified.[18]

[16] Bayne (2010) p. 115
[17] Ibid. p. 197
[18] Bayne (2010) p. 197

Again, the second premise of this argument appeals to a principle outlined by Bayne and discussed above, this time the (CAP). This principle states that all things being equal, phenomenally unified states will be access unified. Given that evidence suggests that few consuming systems will have access to the contents of both 'key' and 'ring' premise (3) seems secure, and with it the argument from access disunity to phenomenal unity.

So, Bayne allows that both the arguments from representational disunity and from access disunity are sound. This gives us reason to think that there is phenomenal disunity in the split-brain cases. I will move on now to address the different conceptions of how this disunity manifests itself, via two-streams and partial unity models of the split-brain cases. Again here I will be drawing on Bayne's (2010) discussion of these issues.

4.4 Two-Streams Accounts of Split-Brain Cases

Two streams approaches to the split-brain cases come (generally) in two types: *contextualists* and *duplicationists*.[19] Contextualists hold that split-brain subjects have two streams of consciousness only under certain experimental conditions, such as those seen in the <key-ring> experiment. Duplicationists by contrast hold that a split-brain subject always has two separate streams.

The contextualist view works in the following way according to Bayne:

> The idea [behind the contextualist view], I take it, is that split-brain patients have what we might call 'scattered experiences'. Consider two neural states, N_1 and N_2, located in the left and right hemispheres respectively. Although each of these states could have realised a K-type experience on its own had the other state been inactive, the patient has only a single token of a K-type experience even when both N_1 and N_2 are active. Rather than N_1 constituting one K-type experience and N_2 constituting another K-type experience, their 'sum' or 'composite' realises a single K-type experience, despite the fact that there are no direct causal connections between them.[20]

Michael Tye is a contextualist with respect to split-brain subjects, and he outlines his reasons for holding this position in his (2003). Tye allows that the split-brain subjects have a single stream of consciousness in everyday situations, but whilst undergoing

[19] These terms are taken from Bayne (2010).
[20] Bayne (2010) p. 201

experiments such as the key-ring experiment outlined above, their consciousness splits. Tye takes the contextualist position to be weaker than the duplicationist line, which has to maintain that a split-brain subject's consciousness is always split, and also takes the contextualist position to be better supported by the evidence of behavioural integration in split-brain subjects in everyday situations. As we have seen, Tye does not hold that a token experience is identical with its neural basis, so the kind of scattered experiences that Bayne takes to be central to the contextualist view are no problem for Tye.

One point which Tye and other contextualists do have to consider is how, in everyday circumstances, a single stream of consciousness can be generated by neural events which are located in two separated hemispheres, and are thus causally unconnected. Tye himself is not unduly concerned by this, and suggests that we should not think that this kind of situation as being so odd.

> It might be objected that a single experience cannot have as its physical basis neural events in the left and right hemispheres that are themselves causally unrelated. But why not? Consider the following example. Two movie projectors each project an image onto a screen at time *t*. Only a single image is present on the screen at t, since identical slides are in the projectors and they are aimed at exactly the same part of the screen. There are two projections but only one image. One projection is redundant. Each projection on its own suffices for the screen image.[21]

The analogy here is with the split-brain subject under experimental conditions, of which Tye says

> There seems no obvious reason why nature should not have made us so that, in certain circumstances, there is redundancy at the neural level in the generation of perceptual experience. After all, it is well known that the human brain has a neurological that is highly redundant anyway. Why not here?[22]

Tye's basis for his defence of this position will come under closer scrutiny in a later section of this thesis.

A further problem for the contextualist two-streams view is how they can explain the shift between everyday and experimental conditions, given that the subject's neural

[21] Tye (2003) p. 127
[22] Ibid. p. 128

structure does not change. These and further objections to the contextualist two-streams model will be addressed in due course. In this section I am simply giving a brief outline of the central tenets of the two-streams views.

In contrast to the contextualist's reliance on scattered experiences, the duplicationist two-streams approach holds that, in the case described by Bayne above, N_1 and N_2 realise two distinct experiences. This is the central commitment of the duplicationist view, and as with contextualism, I will go into more detail concerning the objections later in this chapter.

4.5 Partial Unity Accounts of Split-Brain Cases

The alternative to the contextualist and the duplicationist two streams accounts, if you want to claim that there is phenomenal disunity in the split-brain cases, is to hold that they involve partial phenomenal disunity. Some motivation for this approach comes from evidence of a certain kind of inter-hemispheric integration in the split-brain patients, something the two streams account rules out.

By inter-hemispheric integration here, what is meant is that the patients are capable of integrating *some* information from the two hemispheres. This partial integration is possible due to the fact that different sections of the corpus callosum are responsible for transferring different kinds of information. Tim Bayne details these as follows:

> The anterior mid-body transfers motor information, the posterior mid-body transfers somatosensory information, the isthmus transfers auditory information, and the splenium transfers visual information.[23]

Further to this specialisation of the different sections of the corpus callosum, it has been discovered that it is possible for split-brain patients to exhibit what is known as 'domain-specific splitting'. This domain specific splitting can involve split-brain subjects being split with respect to some senses but not others. For example, a patient (reported by Gazzaniga and LeDoux 1978) who had undergone a partial callosotomy was reported as being split for touch but not for vision. The patient was able to manually retrieve an apple when presented with one on either side of the visual field, with either hand and without using visual exploration as an aid. This is despite tactual information in the left

[23] Bayne (2010) p. 205

hand and right hemisphere being isolated from the right hand and left hemisphere. Patients with full callosotomies can still display some inter-hemispheric integration also. Gazzaniga et al. (1963) describe a patient who was able to locate with his hand, and verbally report, touches applied anywhere on his head and face, despite being unable to locate with the one hand touches applied to the contralateral side of his body on the foot, leg, arm, hand or trunk.

This interhemispheric integration provides some motivation for saying that the phenomenal disunity in split-brain cases may be partial rather than complete. This would involve a single subject with a single but partially unified stream of consciousness. These kinds of cases present a potential problem for the non-reductive views which involve a commitment to transitivity (specifically the views of Dainton, Bayne & Chalmers and Bayne), because of the seemingly partially phenomenally unified stream of consciousness in these cases. A single partially unified stream of consciousness, would, as I have explained, falsify both the transitivity and the unity theses.

Tim Bayne presents a way that this inter-hemispheric integration can be used as an argument for partial unity.[24] He points out that this argument needs to suppose that there are experiences which can 'straddle' the two hemispheres of the subject, but we shall grant this for now. Further, there may not be any definitive reason to rule out these scenarios, at least not a priori. Further still, given what we have seen about the individuation of experiences, demonstrating that such 'straddling' experiences are not legitimate may prove to be a difficult task.

If we consider the example above, of the patient described by Gazzaniga and LeDoux (D.H.) who is split for touch but not vision, Bayne instructs us to imagine a scenario in which D.H. has at one time tactile experiences in both hands, and visual experiences in both hemi-fields. Bayne suggests that in a case like this, we have some reason to think that though the subject's visual experiences are unified with each other, his tactile experiences are not. So, this would make it the case that each of D.H.'s visual experiences is unified with both of his tactile experiences, even though those experiences are not unified with each other. Of this case then, we may say that the subject seems to have a single partially phenomenally unified stream of consciousness. As Bayne says here,

> [I]f that is right, then D.H. has a partially unified consciousness, for (at a single time) he will have a pair of experiences that are unified with each other but not with a third

[24] Bayne (2010) p. 206

> experience… Although D.H. is unusual among split-brain patients in being split for touch but not for vision, we have seen that even the most split of split-brain patients appears to retain the ability to integrate certain types of information, and comparable arguments could be constructed for them also[25]

If we consider this example with respect to the transitivity and unity theses, transitivity fails here because if we take one of the subject's visual experiences, and both the tactile experiences, the visual will be unified with both the tactile, but the two tactile will not be unified. This case would also be a counterexample to the unity thesis, as it would not be the case that all of D.H.'s experiences will be subsumed by a single state, and thus will not be phenomenally unified. There will be a single state which subsumes D.H.'s visual experiences, but not his tactile experiences.[26]

So, this kind of case, and the arguments that can be constructed by considering it gives the partial unity theorist ammunition against those who would claim that there is no phenomenal disunity in the split-brain cases, and those who claim that there cannot be any such thing as phenomenal disunity. Bayne and those others who would advance the transitivity and unity theses theorists then, needs some reason to cast doubt on the partial unity model, before giving an alternative explanation of the split-brain cases. He does this by claiming that although provocative, the argument from integration is not conclusive. To be conclusive, partial unity theorist would need to show that with subjects such as D.H., the relevant behavioural responses are due to inter-hemispheric integration, and this integration is in fact conscious.

An argument for partial unity is also proposed by Michael Lockwood.[27] He argues that it may be possible, in split-brain cases, for the subject's consciousness to fragment gradually, as the corpus callosum was being severed. At some point claims Lockwood, we would see a breakdown in transitivity. This argument would also count against any claims that phenomenal unity is transitive diachronically.

---

[25] Ibid. p. 207

[26] It should be noted here that a case like this one will only involve a failure of transitivity if it is supposed that the phenomenal unity relation is symmetrical.

[27] Lockwood (1989)

4.6 Responses to Split-Brains 1: The Two-Streams & Partial Unity Models Diffused

Before coming to the alternative accounts of what happens in split-brain cases put forward by Bayne & Chalmers, and by Bayne himself, I will look at how the two-streams and partial unity models may be criticised by those wishing to defend the transitivity and unity theses.

I will look at the contextualist two-streams position first. Recall that on this model, split-brain subjects have in everyday conditions a single phenomenally unified stream of consciousness, and it is only under experimental conditions where there is phenomenal disunity in these subjects. Recall also that the contextualist position works via two separate neural events, one in each hemisphere of the subject, jointly realising a single experience, despite there being no causal connection between the two neural states.

We have already seen above that Tye, who is a contextualist, sees no problem with the lack of causal connection between the two neural states and is prepared to embrace some form of neural redundancy in the split-brain contexts. However, it is this lack of causal connection that renders the contextualist view implausible according to Bayne.

> How could the patient's K-type experience produce integrated thought and action if its causal basis is distributed between two causally isolated hemispheres?[28]

Further problems for the contextualist two-streamer are raised by Elizabeth Schechter, who questions the ability of the mereological sum of two distinct neural events to produce a single token experience. This 'singularity through redundancy' principle is relied upon by the contextualist two-streams theorist to account for everyday integration, and Schechter claims it should be rejected.[29] A related problem is how the contextualist two-streamer can explain the shift between everyday and experimental conditions; given that the patient's neural structure does not change. Given these issues, Bayne suggests that we should reject the contextualist two-streams view.

The advocate of transitivity still needs to reply to the duplicationist two-streams view however. The duplicationist holds that rather than that N_1 and N_2 (see above)

[28] Bayne (2010) p. 201
[29] Schechter (2010)

forming the basis of a single experience, the duplicationist holds that N_1 forms the basis for one experience, and N_2 forms the basis for a second, numerically distinct experience.

Bayne's reason for objecting to the duplicationist account is that he claims the idea of duplicated, numerically distinct experiences is dubious and perhaps incoherent. However, the support for this claim is that duplicated experiences do not fit with the tripartite account of individuating experiences which he favours, and which I have discussed above. The tripartite account of experiences individuates them according to subject, time and phenomenal properties, and would rule out duplication in the way the account above requires. Bayne claims we have good reason to keep the tripartite account, and so the duplicationist would seem to be under pressure.

Bayne also give some reasons why any method of individuation the duplicationist could adopt here may be inappropriate. The duplicationist could, as Bayne suggests, individuate experiences in neural or functional terms, so that a single subject might have more than one experience with the same content, as long as those experiences were realised by (or supervene on, or are grounded in) separate neural events or regions. Bayne's objection to this thought seems to be the following: duplication of this kind, along with an account of individuation where experiences are individuated according in neural terms, would mean that these duplicate experiences would not be properties of a subject of experience, but rather of some neuronal assembly or the like. If a subject has two experiences with duplicate content, and what makes these two experiences rather than one is the fact that they involve (in whatever appropriate way) different neural regions, then these experiences cannot properly be said to be experiences which are properties of the subject. Rather, they are properties of the neural regions that realise them. That experiences are states of subjects, not neural regions, should be reflected in their identity conditions, claims Bayne.

While it seems right that duplication and the tripartite account of experience do not sit well together, and may even be mutually exclusive, this may not have the force that Bayne wants, given what we have seen about the individuation of experiences earlier in this chapter and elsewhere in this thesis. I have suggested that when it comes to the question of the individuation of experiences at a time, there are a number of equally legitimate methods of individuation, with no one privileged method. Bayne himself makes the same kind of claim, as we have seen, and suggests there is not any single way in which experiences should be individuated. So, Bayne agrees that there is not any uniquely correct way of individuating experiences. He advocates a tripartite account, as it

fits well with the mereological account of phenomenal unity that he advocates. The duplicationist two streams account of split-brains does not fit well with the tripartite account, but given the lack of a uniquely correct method of individuation, this is not a knock-down objection. Unless Bayne can show that the duplicationist two streamer has to adopt an illegitimate method of individuating experiences (or at least one that is inappropriate in this context, which is itself no small task), then there is not too much pressure on the duplicationist here.

This is a tricky situation, and raises this issue of the synchronic individuation of perceptual experiences again. The dialectic here goes as follows: the duplication of token perceptual experiences proposed by the duplicationist two-streams theorist, and the tripartite conception of the individuation of perceptual experiences which Bayne favours, do not sit easily together. Bayne's motivation for holding the tripartite view comes from its apparent good fit with his mereological account of phenomenal unity. However, Bayne does not hold that the tripartite view is the one correct method of individuating experiences at a time. Indeed, Bayne shares my contention that there may not be any such uniquely correct method. Given this, there seems little pressure on the duplicationist two-streamer to drop duplicate experiences merely on the basis that they clash with the tripartite view.

Further, and the cause of yet more complication, there is again a question of what has priority here, explanation of the split-brain cases and views of phenomenal unity, or the method of experience individuation. As I pointed out when discussing Bayne's view of experience individuation above, a plausible view of phenomenal unity that implies, or relies on, a certain method of individuation will come with pressure to adopt that method of individuation. The opposite however, is also true, in that a plausibly correct or simply favoured view of experience individuation may point to a certain view of phenomenal unity that fits with it. Bayne has a preferred view of individuation, which fits well with his view of phenomenal unity, but as the duplicationist favours another view of unity due to their interpretation of the split-brain cases, then they may be perfectly justified in holding a method of individuation that ties in well with that view. Bayne does not take himself to have refuted the duplicationist view of the split-brain cases, but does take himself to have put significant pressure on it. I am not sure however, that appeal to methods of experience individuation (given the lack of a uniquely correct method for this) can generate such pressure.

We have no knockdown argument against the duplicationist two-streams model of split-brain cases then, and though the contextualist model seems beset by problems, the duplicationist strategy is still on the table.

Concerning partial unity models, despite the appearances of a case such as that involving D.H. described above, Bayne attempts to give us reason to think it possible that the apparent integration in such a case may be due to one hemisphere only. Some doubt about integration comes from the details of these cases themselves: individual hemispheres can sometimes possess a wider range of processing abilities than suspected. These processing abilities include behaviour involving both hemispheres despite a seeming lack of integration between them; inter-hemispheric integration in the absence of consciousness. This kind of empirical evidence, though not conclusive, does provide Bayne with the potential to be deflationary about the kind of seeming integration that leads to partial unity. This inter-hemispheric integration may be explained away by appeal to this evidence.[30]

Even more serious for the partial unity model, claims Bayne, are the conceptual challenges it faces. I will look at two such challenges; that partial unity is *unimaginable*, in the sense that we could not possibly imagine what it would be like to have a single stream of consciousness which was fragmented in this way, and relatedly, that partial unity is in fact *inconceivable*, and therefore impossible.

The argument against partial unity on the grounds of its unimaginability is sometimes referred to as the *projectibility argument*. Since we cannot project ourselves into the mind of a subject whose consciousness is only partially phenomenally unified, we should deny the existence of partial unity. This argument finds some support in Dainton, and in Lockwood (1994) who came to reject the possibility of partial unity, despite his earlier advocacy of it.[31] Bayne offers a reconstruction of the projectibility argument which runs as follows:

(1) If partial unity were possible then there would be something distinctive it is like to be a partially unified subject – there would be such a thing as a partially unified phenomenal perspective.

(2) We are unable to project ourselves into a partially unified phenomenal perspective.

[30] Bayne (2010) p. 209
[31] See Lockwood (1989) and (1994)

(3) If there were such a thing as a partially unified phenomenal perspective then we should be able to project ourselves into it.

(4) Thus, partial unity is impossible.[32]

Note that this argument is not specifically aimed at split-brain cases, but at all putative cases of partial phenomenal unity.

Susan Hurley (1998) counters the projectibility argument by casting doubt on the first premise.[33] Hurley points out that the difference between total and partial phenomenal unity is matter of a relation holding between perceptual states, and since the subject has access only to the content and character of these states and not to the relation which holds between them, then there may not be good reason to accept premise (1) without question.

Though all three of the non-reductive views of phenomenal unity I am currently dealing with also hold that phenomenal unity is a matter of a relation holding between perceptual states, they also hold that the subject has access to the difference made by this relation, and so would expect there to be a subjectively discernible difference between total and partial phenomenal unity.

The premise which is in fact the weakest suggests Bayne, is (3). There is in fact no good reason, he supposes, to take this kind of imaginability (which involves our ability to project ourselves into a given situation) as a good guide to possibility. The projectibility argument fails for this reason. Bayne does however, hope to save something from its wreckage. What he hopes to show is that partial phenomenal unity is in fact *inconceivable*, and therefore, impossible.

This claim seems a natural one for Bayne to make, given that we have seen him to be an adherent to the unity thesis, which states that necessarily, in a single subject at a time, there will be a single state which subsumes all the subject's perceptual states. Since there would not be a single maximal state subsuming all the subject's other states in cases of partial unity, it is by necessity ruled out.

As Bayne himself notes here, inconceivability arguments on the subject of consciousness are decidedly controversial, but still he suggests that it is central to our notion of a phenomenal perspective that phenomenal unity cannot fragment in such a way as to lead to partial phenomenal unity. This claim however, seems purely stipulative on Bayne's part: he has defined phenomenal unity such that it cannot break down, and

[32] Bayne (2010) p. 38
[33] Hurley (1998)

has tried to defend this claim, but to claim that our notion of a phenomenal perspective has this notion of phenomenal unity built into it is also a claim that requires arguing for.

By way of a caveat, Bayne does claim only that partial phenomenal unity is weakly inconceivable, not strongly inconceivable. The difference between theses is that things which are strongly inconceivable, such as square circles, will strike a subject as obviously being so: a square circle is manifestly inconceivable. There are other propositions which are only weakly inconceivable however, such as Goldbach's Conjecture (that every even number greater than 3 is the sum of two prime numbers). The point here is that it is not immediately manifest to us whether the Conjecture is true or false.[34]

Partial unity, suggests Bayne, is weakly inconceivable. Even so, we might question this distinction, and ask whether things which are putatively weakly inconceivable are really inconceivable, or simply not manifestly possible.

Of course, when it comes to conceivability claims, if we are simply left trading intuitions, then neither side makes any progress. One way in which the advocate of partial unity might attempt to demonstrate its conceivability is to take an instance of supposed partial unity, and demonstrate that each of the component parts of that state of affairs are conceivable, and infer from this that the whole thing is conceivable.[35] We can start with a description of the parts of what would be a case of partial unity, and proceed in stages.

(1) Imagine a subject, S, who has at a time three perceptual experiences, e_1, e_2 and e_3.
(2) There is some relation which holds between e_1 & e_2, and between e_2 & e_3.
(3) This relation does not hold between e1 & e3.
(4) Thus there is no one state which subsumes e1, e2 and e3.
(5) There is some conjoint phenomenology to having e1 & e2 together, and having e2 & e3 together.
(6) There is no conjoint phenomenology to having e1 & e3 together.

Taken together, it might be argued that these add up to a case of partial unity. It may also be argued that none of these is inconceivable. Though this strategy may not succeed, it

[34] Bayne (2010) p. 43
[35] I owe this point to Fiona Macpherson.

may at least help highlight exactly what about partial phenomenal unity Bayne takes to be inconceivable by breaking the notion down into its component parts.

I will leave issues of inconceivability however, and move on to look at alternative explanations of the split-brain cases on which there are no breakdowns in phenomenal unity. If any of these alternative explanations are plausible, then the non-reductivist may avail themselves of them to protect the transitivity thesis.

4.7 Response to Split-Brains 2: Bayne & Chalmers' Account

Unlike the two-streams and partial unity models, there are alternative explanations of what is going on in the split-brain cases which do not posit a breakdown in phenomenal unity. In the next section I will look at the 'switch model' currently advocated by Tim Bayne, but before that I will outline a model of split-brain cases developed by Bayne and David Chalmers in their (2003) paper.[36]

Recall that Bayne & Chalmers, like the proponents of the other non-reductive views that I've been examining, hold that phenomenal unity is transitive, because the subsumption relation which they analyse it in terms of is transitive. They claim that there is a prima facie case for the unity thesis being true. They even go so far as to suggest that there may be something incoherent about the suggestion of phenomenal disunity. This tactic is also employed by Bayne, as we have seen above.

Bayne & Chalmers suggest a natural line of response to the split-brain cases, which would preserve phenomenal unity. This response revolves around claiming that in such cases, what is going on is down to a breakdown in access unity, and not phenomenal unity.

The notion of access unity stems from access consciousness, which I introduced in Chapter One. When two or more states are access unified then, as well as being access conscious, they (or their contents) will be jointly accessible.

If we take the following to be our example, we can see what this claim of Bayne & Chalmers' involves. A split-brain subject is presented with two pictures, one on each side of their visual field; a cat on the left hand side and a dog on the right hand side. When the subject is asked to verbally report the contents of their visual field, the subject will only report seeing a dog, since the left hemisphere, which dominates speech, receives its input from the right hand side of the visual field. When asked to write down what they

[36] Bayne & Chalmers (2003)

see, with their left hand (which is controlled by the right hemisphere), the subject will write 'CAT'; with the right hand they may write 'DOG'. The claim that there is a breakdown in access unity goes as follows then: assuming there is a single subject, it seems that in this case they have some weak sort of access to both the presence of the cat, and to the presence of the dog, and can use each in control of reasoning and behaviour, satisfying Block's constraints on a state being access conscious. But despite this, the subject does not have any access to a conjunctive content involving both the cat and the dog, something you and I would have in this situation. No conjunctive content is reported in cases such as this, and plays no apparent role in reasoning and in the control of behaviour.

In order to be germane to our discussion of the transitivity and unity theses, we need to imagine a split-brain case where the subject has three or more perceptual experiences at a time. So, imagine a split-brain subject, who is presented with three pictures; a cat in the left hand side of their visual field, a dog in the right hand side, and a rabbit, intersecting the visual field across the middle, and suppose that when presented with these three stimuli, the subject has three visual experiences. In this case, suppose that the subject would have three experiences; 'CAT' 'DOG' and 'RABBIT' and that 'CAT' would be unified with 'RABBIT' and 'RABBIT' would be unified with 'DOG', but 'CAT' and 'DOG' would not be unified.

A reservation that one may have about this example is that it would be very easy to cast bi-hemispheric representation of the visual midline as a case of two experiences with the same content. In defence of this example, I would point to its simplicity, and to suggest, as I did earlier in reference to experiences which straddle both hemispheres, that we should not rule these things out a priori.

We might naturally think that the disunities in this case, if we think there is disunity, would be in access and phenomenal unity both. Bayne & Chalmers however do not hold that a breakdown in access unity necessarily entails a corresponding breakdown in phenomenal unity. Their defence of this line revolves around their interpretation of the Sperling experiment, and the claim that what happens in it is analogous to what happens in the split-brain cases.[37]

The Sperling experiment involves a subject being presented with a matrix consisting of three rows with four letters in each row. This matrix is flashed up on a screen in front of the subject briefly, for 250 ms, and after the matrix disappears, a tone

[37] Sperling (1960)

sounds, indicating whether the subject is to report the contents of the top, middle or bottom rows, or the contents of the entire matrix. When subjects were required to report the contents of the top row, they correctly reported on average 3.3 of the four letters in the row. The same result was generated when they were asked to report the contents of the middle or bottom rows. But, when the subjects were asked to report the contents of the entire matrix, they only managed on average to report 4.5 out of the twelve letters correctly.

The results of the Sperling test suggest essentially that the subject can report any four letters on command, but cannot report any combination of over four letters. So, in light of these results, Bayne and Chalmers suggest that the subject in this experiment has access to the information in each single row, but no joint access to all three rows. It is natural to suppose, they suggest, that just after the matrix disappears, and just before the tone sounds, the subject is access-conscious of the contents of each individual row, but is not access-conscious of the conjunctive content of all three rows.[38]

This initially seems a reasonable interpretation of what is going on in the Sperling case, nevertheless it is important to see what exactly a breakdown in access unity involves in the Sperling experiment[39]. The breakdown here would be one of joint accessibility: of the three rows of letters, the first row is accessible on its own, as are the second and third rows, however, row 1 and 2, or 2 and 3, or 1 and 3, are not *jointly* accessible. This is demonstrated in the following table.

Table 4.1. Access Breakdowns in the Sperling Experiment

Rows	Accessible?
1	Yes
2	Yes
3	Yes
1 + 2	No
1 + 3	No

[38] Here I am keen to avoid controversy over what exactly the subject has access to; mental states, or the items in the world. I believe there is an intuitive sense in which it makes sense to talk of the subject's having access to the contents of the rows of letters, without getting bogged-down in controversy.
[39] For an alternative analysis of Sperling, see Phillips (forthcoming) and (forthcoming) (b)

2 + 3	No
1 + 2 + 3	No

As the table above shows, what is breaking down in the Sperling Experiment is joint accessibility: any *one row* is accessible to the subject, but any *combination of rows* is not.

Access unity can break down then, but according to Bayne & Chalmers, this fact need not entail anything about breakdowns in *phenomenal* unity. They suggest that there is little reason to think that there is not phenomenal unity in the Sperling Experiment:

> No matter what it is like for a subject to experience each individual cell of the matrix in the Sperling case, it is plausible that there will be something it is like for the subject to see the entire matrix. And it is plausible that the phenomenology of seeing the matrix will subsume the phenomenology of seeing the individual cells. If the phenomenology of seeing a cell involves just a hazy patch, then the phenomenology of seeing the matrix will plausibly involve nine hazy patches. If the phenomenology of seeing a cell involves a detailed shape, then the phenomenology of seeing the matrix will plausibly involve nine detailed shapes. Either way, the individual phenomenal states are subsumed by the overall phenomenal state. So there is no reason to deny phenomenal unity here.[40]

Bayne & Chalmers take it as important evidence in their favour that their claim is also backed up by the subject's reports when they undergo the experiment; they claim that there is something it is like for them to see the whole matrix of letters before it disappears.

So, if Bayne & Chalmers are right, and if a breakdown in access unity need not entail any breakdown in phenomenal unity (as shown by the Sperling case) then the possibility remains that in split-brain cases also, there is no failure of transitivity, as these partial split-brain cases may also involve a failure of access unity without any corresponding failure of phenomenal unity.

As an aside, I should point out again here that there is a huge debate on what precisely the relationship between access and phenomenal consciousness is. Are these just two ways of thinking about the same phenomenon, or are they in fact distinct phenomena? If they are distinct phenomena, can they come apart? There are those who have argued that the best scientific assumption to make is that all phenomenally

[40] Bayne & Chalmers (2003) p. 36

conscious states are also access conscious, since, the claim goes, as a methodological point the only evidence there could be for a state's being phenomenally conscious is through people's reports of their experience. And so a similar claim could be made about access and phenomenal *unity*. According to this claim, we would only really have evidence of access disunity in split-brain cases and not of phenomenal disunity. There are however objections to both these claims: it is pointed out that this strategy risks losing the evidence we have for the existence of phenomenal consciousness in the first place, and that this strategy is also guilty of cherry-picking when to use introspective reports of experiences. Another issue concerns our lack of a test for phenomenal consciousness, other than relying on verbal reports (i.e. relying on the presence of access consciousness). We may worry about what kind of test we could currently employ to demonstrate that a given case was one of phenomenal consciousness (or indeed unity) without access consciousness. I do not have the space to pursue these issues any further here however, and am merely flagging up their presence. A further point to note here is that Bayne & Chalmers are relying on what has become the received view of the Sperling experiments, that phenomenal consciousness outstrips access consciousness in these cases. There are alternative interpretations put forward however, by people such as Ian Phillips.[41]

Returning to the matter at hand then, though this strategy of Bayne & Chalmers' seems initially to be a promising response that the likes of Dainton may also apply to the split-brain cases, in actual fact this line of response has problems which I think make it impossible to use for the purpose of successfully defending the transitivity thesis.

The problems for the Bayne & Chalmers line stems from the fact that there is actually a disanalogy between the type of breakdown in access unity that goes on in the Sperling cases, and the split-brain cases. If access unity breaks down in the split-brain cases, then it would break down in a way which is *not* analogous to the access unity breakdown that occurs in the Sperling Experiment.

The kind of access unity breakdown that would occur in the split-brain cases like the one outlined above would involve the subject having conjoint access to 'CAT' & 'RABBIT' and conjoint access to 'RABBIT' & 'DOG', but no conjoint access to 'CAT' & 'DOG'. This kind of access breakdown is illustrated in Figure 4.2 below, and a comparison between the Sperling and split-brain cases shown in 4.3.[42]

[41] See Phillips (forthcoming) and Phillips (forthcoming) (b).
[42] Here 'C' stands for 'CAT', 'R' for 'RABBIT' and 'D' for 'DOG'.

Table 4.2 Access Unity Breakdown in Split-Brain Cases

Experiences	Accessible?
C	Yes
R	Yes
D	Yes
C + R	Yes
C + D	No
R + D	Yes
C + R + D	No

Table 4.3 Access Breakdowns in the Sperling and Split-Brain Cases, Compared

Sperling		**Split-brain**	
Experiences	Accessible?	Experiences	Accessible?
1	Yes	C	Yes
2	Yes	R	Yes
3	Yes	D	Yes
1+2	No	C+R	Yes
1+ 3	No	C+D	No
2+3	No	R+D	Yes
1+2+3	No	C+R+D	No

From the tables above, we can see that the breakdown in access unity that would arise in the split-brain cases, is not the same as the breakdown in access unity that occurs

in the Sperling experiment. There are two crucial differences between the kind of access unity breakdown that appears in the Sperling experiment and those that appear in split-brain cases. These differences between the two kinds of access breakdown can be further explained as follows:

(1) The first difference is that what can be jointly accessed in the Sperling case is *limited solely by number*. The subject can report any four letters, but no more than that, whereas in the partial split-brain case, the limit is not solely numerical; *some* combinations of two stimuli can be jointly accessed, while other combinations of two stimuli cannot. While this in itself does not constitute a reason to think that partial unity of access consciousness is not present in the split brain cases, it does provide a reason not to draw the analogy with the Sperling case.

(2) The second difference is that in the Sperling experiment, any single stimulus that is actually accessed, can be accessed jointly with any other stimulus that is accessed. In the partial split-brain case however, this is not the case. Here, though three stimuli may all actually be accessed individually, so doing does not ensure joint access to any two of the three, or to all three.

To elaborate on the first difference between the two cases still further: even though initially the split-brain cases and the Sperling cases are the same with respect to the joint accessibility of rows 1+3 in the Sperling case, and experiences 'CAT' and 'DOG' in the split-brain case, it would be possible to make these two results diverge, whilst retaining the set-up of the Sperling experiment. So, if each row was made to contain only two letters instead of three or four, then it may appear on at least some trials that rows 1+2 were jointly accessible, and on some trials 2+3 were, and on some trials 1+3 were jointly accessible. This result is in principle possible given what is going on in the Sperling experiment, however experiences 'CAT' and 'DOG' will never be jointly accessible in the split-brain cases, as here the limits of joint accessibility are not limits solely of number, but limits dictated by the absence of a physical pathway that could enable the joint access of the two experiences.

Why does the difference between the two types of access breakdown make a difference to whether or not they could accompany phenomenal unity? Explaining why this is so is tricky, and requires some work to provide an answer.

Firstly, I will try to motivate the thought that in the Sperling case, it is far more natural to think that we could have a breakdown in access unity without a breakdown in phenomenal unity than it is in the partial split-brain cases. In the Sperling cases, the reason why all three rows of the matrix are not jointly accessible is plausibly due to the natural limits on our ability to jointly access information: under the conditions of the experiment, a single row is all we can jointly access. Here we needn't posit a breakdown in phenomenal unity to explain what is going on in these experiments, what is going on is simply that there is a natural limit to our capacities to jointly access information. This is not what is going on in the split-brain cases. Indeed, it may even be more accurate not to describe the Sperling case as involving a genuine *breakdown* in access unity at all, but rather a processing bottleneck. In his more recent treatment of the split-brain cases, Bayne himself acknowledges this. Bayne also concedes here that the analogy between the split-brain cases and the Sperling experiment breaks down with respect to the *kind* of access the subject has to the contents of their experience.

> Furthermore, where there are cognitive bottlenecks, we are usually able to report that we are aware of more than we can directly report. Subjects in the Sperling experiments cannot report the contents of their experience of the entire matrix, but they can (and do) indicate *that* they had an experience of the matrix whose contents outstripped what they could report. But S [a split-brain subject] fails to have even this indirect form of access to her experiential content.[43]

The second difference between the two cases is that in the Sperling case, actual access of any one of the stimuli guarantees joint access to the rest of the stimuli, up to the subject's limit for joint access. So, when the subject accesses one of the letters in a row, this means that joint access to the rest of the row is ensured. Again, this does not happen in the split-brain case, where though all three of the stimuli may be accessed individually, this does not ensure joint access to all three, or even two. This second difference provides us with more reason to think that in the partial split-brain case, the kind of access breakdown that occurs could not occur whilst phenomenal unity remained intact: we may think that in the Sperling case, access of one stimulus all but ensures joint access of the rest, *because* of the presence of phenomenal unity. In the partial split-brain cases, the strange kind of access breakdown that we would see is best explained by a

[43] Bayne (2010) p. 198

breakdown in phenomenal unity.[44] So, in the Sperling cases, we can explain the breakdown in joint accessibility without appealing to a breakdown in phenomenal unity, and there are other reasons to think that in fact phenomenal unity is present, but in the partial split-brain cases, the best explanation of the kind of access breakdown that we would see here may be that phenomenal unity has broken down also.

So, Bayne & Chalmers have given a story suggesting why we should suppose that a certain kind of access unity breakdown (illustrated by the Sperling experiment) is always accompanied by a corresponding breakdown in phenomenal unity. They have not shown however that we can have a breakdown in access unity, of the kind that would occur in the split-brain cases, without an accompanying breakdown in phenomenal unity. Their position only goes so far as to suggest that the kind of access breakdowns we see in the Sperling experiment can nevertheless accompany phenomenal unity. They show nothing about phenomenal unity being present alongside the kind of access unity breakdowns that would be present in the split-brain cases. Thus the unity and transitivity theses are still not secured.

4.8 Response to Split-Brains 3: Bayne's Switch Model

Given that we have seen above the flaws in Bayne & Chalmers' attempt to diffuse the split-brain cases, Bayne needs another approach in order to defend the unity and transitivity theses in his more recent work. His alternative account of the split-brain cases which seeks to do just that is called the *switch model*. The basis of the switch model is the idea that rather than supposing the split-brain patient's two hemispheres are conscious in parallel, with or without integration between them, we should hold that consciousness switches from one hemisphere to the other. Though both hemispheres can process information concurrently, they alternate when it comes to supporting consciousness. so, in the key-ring experiment, the patient might be conscious of the word 'key' due to activity in the right hemisphere, or may be conscious of the word 'ring' due to left hemisphere activity, but since consciousness is never supported by both hemispheres simultaneously, the patient will never be conscious of both 'key' and 'ring', even when they are simultaneously presented to them.

Firstly, what reason do we have to choose Bayne's switch model over the other interpretations of the split-brain cases?

[44] Tye (2003) makes this same point, in defence of the two-streams model.

Bayne cites as important evidence for the switch model studies involving chimeric stimuli, which are stimuli that are created by joining two similar half-stimuli at the patient's vertical midline of their visual field. Such chimeric stimuli can be created from faces, using the left side of one face, and the right side of another. In experiments with these stimuli, the subjects were asked to carry out tasks such as pointing to a matching stimulus, or naming the stimulus. On almost all of these trials, the patient indicated only one match for each of the chimeric stimuli.

What is happening in these experiments? Bayne cites the account of Jerre Levy, which is captured by the following passage, worth quoting in full:

> With one half-stimulus joined at the mid-line to a different half-stimulus to make a 'chimera', each hemisphere would receive equivalent, but different stimulus input, and if two perceptions were gained, they would be in conflict as evidence concerning the object of interest, and the motor responses guided by these percepts would be in conflict… [But] if… perception is the preparation to respond, then, except in special circumstances, there should be a single perception linked to a single response under conditions of competitive stimulus input… Our studies overwhelmingly confirmed the predictions of this conceptual model. For all patients examined, and for tasks including the perception of faces, nonsense shapes, pictures of common objects, patterns of Xs and squares, words, word meanings, phonetic images of rhyming pictures, and outline drawings to be matched to colours, patients gave one response on the vast majority of competitive trials. Further, the nonresponding hemisphere gave no evidence that it had any perception at all. Thus, if the right hemisphere responded there was no indication, by words or facial expression, that the left hemisphere had any argument with the choice made, and, similarly, if the left hemisphere responded, no behaviour on the part of the patient suggested a disagreement by the right hemisphere.[45]

So, from these experiments and their results, it is concluded that subjects were reporting based on information processed by one hemisphere, which was the one that supported consciousness at that time. Bayne describes the patients as undergoing a kind of 'fluctuating perceptual extinction'.

This seeming extinction leads to the view that when one hemisphere is active, say the right, stimuli presented to the left visual field are ignored in favour of those presented

[45] Levy (1990) quoted in Bayne (2010) p. 211

to the right, and further, the lack of any response from the other hemisphere suggests that the right hemisphere is the only one supporting consciousness at that time.

Bayne also points to a follow-up study, related to Levy's work, where all but one of the test subjects involved exhibited the same kind of extinction, and reported only those stimuli restricted to a single hemi-field, and a further study in which perceptual extinction was seen in a patient who was tested for the bilateral integration of tactile stimuli.[46]

Bayne takes there to be further support for the switch model, coming from studies of auditory processing in split-brain patients. Though information from the ear is sent to both hemispheres, there is not the same kind of lateralisation for auditory information in the split-brain patients, as there is lateralisation of visual and tactile information. Nevertheless, lateralisation can be induced via an experiment where subjects are simultaneously presented with competing auditory stimuli to each ear. Under these conditions, the competition between the hemispheres means that seemingly only the information from one ear enters consciousness at any one time, and which hemisphere dominates can be adjusted depending on the task the subjects were asked to perform.[47]

From these experiments, Bayne draws the conclusion that here also, the split-brain subjects' consciousness is switching, from one hemisphere to the other, and that in these cases, the switching can be modulated by the demands of the task set in the experimental conditions. This leads him to revise the status of the arguments from access and representational disunity set out above. These arguments are the standard ways to argue for phenomenal disunity in the split-brain cases, but rely on the split-brain subject's two hemispheres being simultaneously conscious. The switch model however, entails that the two hemispheres are not simultaneously conscious, and so if correct, the switch model would rule out arguing for phenomenal disunity in the split-brain cases on the basis of access or representational disunity.

> Consider again the key-ring experiment. The representational and access disunity arguments assume that the patient's two hemispheres must be simultaneously conscious because the stimuli are simultaneously presented to the patient's visual hemifields, and because each hemisphere can respond to the stimulus its hemifield as and when required. But both inferences are contentious. Perhaps the ability of patients to respond in this way is the result of consciousness switching rapidly and effortlessly between hemispheres in response to the demands of the patient's context. The hemisphere that is silent on any

[46] See Teng & Sperry (1973) and Gazzaniga et al (1963)
[47] See Milner et al (1968) and Milner et al. (1990)

> one trial may be so because it is unconscious rather than because it is unable (or unwilling) to 'speak'.[48]

So, Bayne takes himself to have explained how the switch model can account for the experimental data from the split-brain cases. However, he has so far only considered evidence from studies that seem to back up the switch model itself. What of the experimental data that would be cited by partial unity theorists? Also, how does the switch model account for the everyday integration of split-brain patients?

To take the latter question first, Bayne outlines three approaches which the switch model theorist might take. Firstly, it's possible that split-brain subjects get through everyday life using only a single conscious hemisphere. Secondly, it may be that the interhemispheric switches that occur according to the switch model are so smooth and rapid that they generate the impression that the subject is conscious of much more then she happens to be. Lastly, it may be that for much of everyday life, the split-brain subjects do not require the kind of focussed attention that is required under experimental conditions. Bayne himself does not commit to any one of these options, but suggests that there is enough promise in them to mean that the switch account is a viable one.

What of other objections to the switch model? How does Bayne respond to the data that previously seemed to support a partial unity model? This data suggested that in experimental conditions, the split-brain patients had experiences which straddled their two hemispheres, leading to a failure of transitivity and of phenomenal unity. If Bayne's switch model is to succeed in the task that he has set it, it needs to be able to diffuse the partial unity claims.

If split-brain subjects do have experiences which straddle both their hemispheres, is this necessarily at odds with the switch model? Bayne claims that it is not. The crucial issue is whether or not the straddling experiences are unified with other experiences which are not also unified with each other. This situation I have suggested is possible, and it puts serious pressure on any primitive relation theorist who would hold that phenomenal unity is transitive and holds necessarily within a single subject, as both Dainton and Bayne do. Consider again a split-brain patient who is split for touch but not for vision. This patient is presented with visual stimuli to both hemifields and tactile stimuli to both hands. Previously, we have seen that this case may lead to partial phenomenal unity, but Bayne denies the possibility of partial phenomenal unity, and so

[48] Bayne (2010) p. 213

needs a way of reinterpreting this case, using the switch model, so that partial unity is avoided. Perhaps it is possible here, he suggests, that the subject's visual experiences can be unified with one tactile stimulus or the other, but not with both at once. A more general possibility that Bayne countenances is that while consciousness can straddle different hemispheres at different times, the kind of scenario required for partial unity is not possible. This Bayne concedes is speculative, and relies on our being able to know much more about what subjects such as this can be simultaneously conscious of. This task itself is complicated further by the issues surrounding access and phenomenal consciousness which I have previously brought up.

Thus, Bayne cannot outright reject the possibility of partial unity, even with his switch model in place. I will return to this point shortly, after taking a brief look at one further objection to the switch model.

The objection is concerned with why the split-brain patients do not report sudden changes in the contents of their consciousness. The switch model holds that the contents of the subject's consciousness are informed by what is going on in each hemisphere, and so should change when the switch between hemispheres takes place. However, the subjects do not report such sudden changes in the contents of their consciousness.

This Bayne takes to be the most serious objection to the switch model. In response, he appeals to extensive representational overlap between the contents produced by both hemispheres, such that the subject's perspective on the world does not shift radically, depending on which hemisphere is primarily informing consciousness. However, there will of course be occasions when the shift between hemispheres *does* produce substantive shift in the contents of consciousness. Even here however, Bayne suggests that we might have reason to think that the subjects will not report this shift.

The reasons here stem from the fact that there is some difference between changes in the contents of consciousness, and conscious representation of those changes. There are various impairments which can prevent subjects from noticing changes to the contents of their own consciousness. Subjects with achromatopsia (agnosia for colour) are frequently unaware of the change brought about by losing their colour vision. Subjects who have unilateral neglect are also disposed not to notice gaps or changes in the contents of their consciousness. Change blindness experiments have demonstrated that even subjects without neurophysiological impairment can fail to notice changes in the contents of their consciousness. Finally, Bayne notes that perhaps it is possible that introspective access to an experience requires the neural areas that generated that

experience to be active. If this were so, then in order to be aware of switches in the contents of their consciousness, the subjects would have to have to simultaneously activate the relevant areas in both hemispheres, something the switch model rules out. All of this means Bayne seems able to save the switch model from this objection. However, given that the switch model does not rule out partial phenomenal unity in the split-brain cases, then though it may be a plausible model of what happens in some split-brain subjects, it does not do the work that Bayne needs it to do.

The reason that studies of the unity of consciousness focus so frequently on the split-brain cases is that they seem initially to be good candidates for cases where the unity of consciousness breaks down in some way. As we have seen from Dainton, Bayne & Chalmers, and now Bayne, some people develop theories of phenomenal unity which entail that it cannot break down. So, for these people, the split-brain cases need to be diffused or explained away in some way.

Though Bayne defends the switch model of split-brains, if this model does not allow him to rule out any kind of phenomenal unity breakdown, then it does not help in defending the unity thesis that he holds, whereby phenomenal unity cannot break down. Given that Bayne concedes that the switch model cannot rule out partial phenomenal unity, then the unity and transitivity theses are not secured.

Further, as we have seen in previous sections of this chapter, the attempts to diffuse the two-streams and partial unity models of split-brains have not been conclusive, and Bayne & Chalmers' method for preserving phenomenal unity in these cases does not succeed. So, if Bayne's switch model cannot rule out phenomenal unity breakdowns in the split-brain cases, then as things stand at the close of this chapter, the unity and transitivity theses held by Bayne, Bayne & Chalmers and by Dainton, are under threat.

4.9 Conclusion

What is common to all three of the accounts of phenomenal unity looked at (though as I have noted, not necessarily common to all non-reductive accounts), as well as treating phenomenal unity as something to be explained without reducing it to some other kind of unity, is all three view's commitment to phenomenal unity being a transitive relation.

As I pointed out when discussing the non-reductive views, there are some major benefits to a theory on which phenomenal unity is transitive: above all, it provides a very neat picture of a subject's stream of consciousness at a time. However, after examining

the responses to the split-brain cases made by these non-reductive views, it seems that it is this commitment to transitivity which should lead us to reject the non-reductive views of phenomenal unity advanced by Bayne, Dainton and Bayne & Chalmers.

From the discussion of the split-brain cases which I have investigated in this chapter, there seems nothing especially conclusive at this stage, which would force us to favour one account of the split-brain cases over another. Though this does not in itself damage the transitivity and unity theses, it does mean that the advocates of these theses have not yet done enough to rule out breakdowns of phenomenal unity in the split-brain cases. Further, Bayne himself admits that the pursuit of *the* one structure of split-brain cases may be jeopardised by differences between split-brain subjects, and even within a subject. Given this, it seems that partial unity in the split-brain cases has not been ruled out empirically, and despite Bayne's claims, may not be ruled out conceptually either. This puts the three accounts of phenomenal unity examined in this section under pressure, as they have not ruled out the counterexamples to their transitivity claims.

We would be in a happier position if we did not have to impose a blanket theory on all split-brain cases then, and we could do so if we were not under pressure to preserve the unity thesis, or maintain claims of transitivity. This should serve as warning against our adopting the non-reductive views of phenomenal unity which make explicit commitment to either the transitivity or the unity theses. This itself does not mean that there could be no workable non-reductive explanation of phenomenal unity, one which did not share the commitments of the extant theories I have looked at (indeed, I will go on to advocate just such a theory). A non-reductive view could be outlined which is simply silent about the logical properties of the phenomenal unity relation, and treats it simply as a primitive about which we can say nothing more. A view such as this however will look explanatorily unappealing next to a reductive view which can give a much fuller explanation of phenomenal unity, and can also avoid the need to impose a blanket model of unity on the split-brain cases and other similar cases.

This concludes my discussion of non-reductive explanations of phenomenal unity. In the next section I will move on to look at reductive explanations, where phenomenal unity is explained in terms of some other kind of unity exhibited by consciousness. I will examine and subsequently reject two such accounts, before advancing my own favoured explanation of phenomenal unity, which as we shall see, ends up being a non-reductive account in the same sense as Bayne and Bayne & Chalmers' accounts, as it will not seek to reduce phenomenal unity to any other kind of unity.

Section Three – Reductive Accounts of Phenomenal Unity

Chapter Five - Phenomenal Unity as Spatial Unity

5.1. Introduction

In the previous chapter, I gave my reasons for not adopting a non-reductive account of phenomenal unity. In this section of the thesis, I will look at two reductive accounts of phenomenal unity which share the feature that they explain phenomenal unity solely in terms of another kind of unity exhibited by consciousness. On these reductive explanations, the conjoint phenomenology associated with phenomenal unity comes in fact from some other kind of unity. The first of these such reductive accounts seeks to explain phenomenal unity in terms of spatial unity, and so two or more perceptual experiences are phenomenally unified when and only when they are spatially unified.

The form of the chapter will be as follows: firstly I will provide some background to why one might think that it is possible and indeed plausible, to give a reductive explanation of phenomenal unity that reduces phenomenal unity entirely to some other kind of unity that consciousness exhibits. Next I will suggest why a reductive explanation of phenomenal unity in terms of spatial unity in particular may be thought to be plausible, suggesting that there is intuitively a lot to recommend this explanation. I will then move on to providing some further description and explanation of what it is for two or more perceptual experiences to be spatially unified, and further flesh out the idea of a single common space. I will then look at the reasons why Barry Dainton rejects an explanation of phenomenal unity in terms of spatial unity. Though ultimately, I think that Dainton is right to reject such an explanation, there are problems both with Dainton's arguments against the spatial unity thesis, and his characterisations of spatial unity itself. However, we can modify the arguments Dainton gives in order to show that spatial unity is neither necessary nor sufficient for phenomenal unity, and I will do so in this chapter.

Thus the conclusion of this chapter will be that an attempt to reductively explain phenomenal unity in terms of spatial unity fails.

5.2. A Reductive Explanation of Phenomenal Unity?

So, why should we think that it is plausible, or even possible, to explain phenomenal unity solely in terms of another kind of unity that consciousness, exhibits? Before addressing this question however, I will explain what I am taking a reductive explanation to be in this thesis and why at certain points I use the term in a non-standard way.

We have already seen that exactly which views of phenomenal unity should be classed as reductive ones is not immediately clear: I have grouped together the views of Tim Bayne, Bayne & Chalmers, and Barry Dainton under the heading 'non-reductive' views, but as I have already explained, this grouping is supposed to indicate that on all three of these views, phenomenal unity is not reduced to some other kind of unity. In the strictest sense, the only non-reductive view of phenomenal unity is Dainton's, on which (as we have seen) phenomenal unity is taken to be a primitive, sui generis relation, about which nothing further can be said, by way of explanation. In grouping these views together in this way, and in calling them non-reductive views of phenomenal unity, I am using 'reductive' and 'non-reductive' in a slightly non-standard way, but this is purely for dialectical purposes, to highlight the differences between the views laid out in the previous section, and those which will be examined in this section of the thesis.

In this section of this thesis however, I will be examining explanations of phenomenal unity which do seek to reduce phenomenal unity to some other kind of unity that consciousness exhibits, be it spatial or introspective unity. The contrast between the views discussed in this section and those discussed in the previous section is best described as being that between views which seek to reduce phenomenal unity to some other kind of unity, and those (discussed in the previous section) which do not. Importantly however, when in either case I make a claim about an explanation of phenomenal unity being reductive, the notion of reduction itself I am appealing to is the same, involving necessity, sufficiency, and 'nothing over and above' claims.

Before moving on to examine the potential of reducing phenomenal unity to spatial unity specifically, I will briefly make some more general comments concerning reductive explanations.

The motivations for attempting a reductive analysis of some phenomenon that stands in need of explanation should be reasonably clear and uncontroversial: it seems generally prudent to attempt reductive analysis, and where possible, avoid positing primitives. As a more specific point, we have also seen the problems associated with those views which do not seek to reduce phenomenal unity to some other kind of unity, so views which do make such a reduction are worth consideration.

Now, a further issue that merits comment at this stage concerns what exactly a reductive explanation is, and what it means to say that phenomenal unity, which is a relation which holds between concrete particulars, is reduced to something else.

This in itself is a controversial issue, with there being no great consensus about the mark of a reductive explanation. Addressing this issue in detail is outwith the scope of this thesis, so I will restrict myself here to what I take a reductive explanation to be, with respect to the theories examined in this section of the thesis. In saying that, for example, phenomenal unity reduces to spatial unity, I will be taking this to mean that spatial unity is both necessary and sufficient for phenomenal unity, and involving a 'nothing over and above' claim, where there is no feature of phenomenal unity that is not possessed by spatial unity also. In other words, phenomenal unity is exhausted by spatial unity, and thus we can reduce the former to the latter.

To return to the main focus of this chapter, the idea that phenomenal unity could be reductively explained in terms of spatial unity has a certain intuitive plausibility on first considering it.. It seems that in the case of perceptual experience at least, that it is inherently spatial inasmuch as our perceptual experiences always have some spatial element, and that something like this could be responsible for the phenomenal unity of perceptual experiences.

To illustrate the above point, consider some everyday perceptual experiences, you are sitting at your desk watching the rain on the window. At the same time, you can hear music from the radio, smell coffee, and feel the chair against your back. You also have proprioceptive experience of the position and movement of your body. In thinking about this fairly mundane example, notice that there is a common spatial element to all your perceptual experiences here: you see that the rain hitting the window is a certain distance in front of you; the music from the radio is coming from over on your left; the coffee smell from the cup down on your right, and the pressure of the chair on your back is felt to be located relative to you, and to the other objects around you. Your overall perceptual experience at this time then, contains various spatial elements and relations.

The objects of your various experiences are represented as being located somewhere in space, both relative to you, and to each other. The objects of your perceptual experiences here (including your own body) are all represented as being located in the same common space.

Perhaps this spatial element of perceptual experience can also be responsible for the conjoint phenomenology of phenomenal unity. Perhaps there is something it is like for you to experience the rain lashing against the window, and hear the radio together, *because* of the spatial connectedness of your perceptual experiences. This may offer an easy response to the split-brain cases also: the reason why the split-brain subject does not enjoy any conjoint phenomenology when presented with 'key' and 'ring' is because they are not spatially connected in the subject's visual experience. Further, a view on which phenomenal unity was explained by spatial unity would seem to be well placed to explain how conjoint phenomenology involves extra phenomenology, if indeed it does. Since the difference between phenomenally unified and disunified states would on this view be the difference between spatially unified or disunified states, the potential extra phenomenology would come from the objects of experience being represented as being located in a single common space. Even if conjoint phenomenology does not equal extra phenomenology (I am still remaining neutral on this) spatial unity seems initially to be well placed to explain this.

As an important note here, as I am not dealing in this thesis with the phenomenal unity of states other than perceptual experiences, I make no claims about the spatial content of other states. Further, it will not count as a counterexample to the thesis I am considering in this chapter if it turns out that desires have no spatial content, for example.

Before I move on to address this potential explanation of phenomenal unity in detail however, I will first go into some further detail on some key points concerning spatial unity itself.

5.3 Spatial Unity

As we saw in Chapter One, spatial unity is a form of *representational unity*. As I said previously, representational unities concern not the bearer, or subject of states or experiences, but the objects or contents of these experiences or states. States are representationally unified to the extent that their contents are integrated with one

another. Spatial unity then concerns integration of perceptual objects in a single space, and thus two states are spatially unified iff the objects of those experiences are represented as being located in a single common space. The first and simplest way in which two experiences could be said to be spatially unified is if the objects of the two experiences are represented as being located in the same objective space. What is meant by *objective* space? In the most basic terms, if the objects of an experience are represented as being located in objective space, then they are represented as being located in the world, independently of any one subject's perspective or point of view. This kind of spatial unity posits only very slight demands on the spatial content of perception, as experiences need only represent their objects as being located in space. Nevertheless, this kind of spatial content is not vacuous, as there are potentially perceptual experiences which do not represent their objects as being spatially located at all.

The second way in which two experiences could be spatially unified is when the objects of those experiences are represented as being located in the same *egocentric space.* Egocentric space is a widely used term in contemporary philosophy of mind, but it is worth taking some time to outline the notion further.

Egocentric space is a spatial framework which has its axis centred on the subject's body. It is traditionally thought that the centre of the axes is the subject's torso rather than the head, and I will take this traditional conception of the nature of the egocentric framework to be correct. There are of course other views about how a subject's egocentric space is structured, such as the idea that a subject has a multiplicity of non-competing egocentric spaces, demarcated according to limbs/body parts and their possibilities for movement. It is beyond the scope of this thesis however, to conduct a detailed investigation into the different accounts of egocentric space, and so I will stick with the traditional conception. Different accounts of egocentric space may affect the spatial unity explanation of phenomenal unity of course, and so I will be careful to flag up when this may be so.[1]

5.4 Phenomenal Unity as Spatial Unity?

[1] Discussions of egocentric space can be found in Campbell (1994) Evans (1982) and Brewer (1995) amongst others.

Having spelled out spatial unity in some more detail, we can now continue with the idea that phenomenal unity can be explained in terms of some kind of spatial unity. Barry Dainton discusses this thesis, and he outlines it as follows

> [S]imultaneous experiences are co-conscious solely by virtue of occurring at the same time within a single unified three-dimensional space; being thus spatially connected is both sufficient and necessary for co-consciousness.[2]

Dainton refers to this as the *S-thesis*. Though here Dainton casts this S-thesis in terms of *experiences occurring in the same space*, it seems that the thesis would be far better cast in terms of the *objects of experience* being represented as being located in the same space. This does the job just as well, and it steers the S-thesis away from the contentious issue of where it is that experiences take place.[3] In fact, Dainton himself carries on much of the discussion as though the S-thesis were formulated as I have suggested. The S-thesis should run as follows then (I will swap Dainton's co-conscious nomenclature for my own):

> Two or more simultaneous perceptual experiences had by a single subject are phenomenally unified when and only when they are spatially unified.

We can plug into this S-thesis any of the versions of spatial unity, with the variable here being the conception of a single common space that is involved. I will go on to show however that despite the initial plausibility of the S-thesis described above, spatial unity is in actual fact neither necessary nor sufficient for phenomenal unity.

As an important caveat here, it should be noted that spatial unity does not require access unity. That is, two experiences had by a subject at a time need not be such that the subject has any conjoint access to the spatial content of the two states. Two experiences could be spatially unified, in that their represented objects are represented as being located in the same common space, without their being access unified, i.e. without the subject having conjoint access to the contents of the two experiences. This point will re-emerge later in the chapter.

---

[2] Dainton (2006) p. 61 Recall I am taking 'co-conscious' and 'phenomenally unified' to be synonymous in the context of this thesis.

[3] This move also avoids commitment to the sense-datum framework Dainton himself tacitly seems to adopt.

In showing that spatial unity is neither necessary nor sufficient for phenomenal unity, I will draw on Dainton's discussion of this issue, but although Dainton and I are ultimately agreed that the S-thesis fails, there are I think flaws in Dainton's arguments, and further, problems with the way that Dainton introduces the idea of spatial unity. Firstly, as we have seen, Dainton talks in terms of the experiences themselves being located in the same space, which is potentially confusing, and out of sync with how spatial unity is characterised by others working in this area. Secondly, in cashing out spatial unity, Dainton talks not of objective, or egocentric space, but of *phenomenal space*. What a phenomenal space is, is not immediately perspicuous, so I will try and explain it here. As a first pass, a phenomenal space as Dainton describes it is a single three-dimensional field, which possesses a certain dimensionality and size, egocentric axes, and further, possesses its own distinct phenomenal character.[4] The size of the phenomenal space need not necessarily be constant, and it and the structure of the space will be fixed by the characteristics of the corresponding sense field. What it means to say that a phenomenal space has a size and structure that is set by the corresponding sense-field, a size which need not be constant is to say that the size and structure of a phenomenal space will correspond roughly to the size of the visual field, though it will fluctuate in size due to the input of the other senses. For instance, if I hear a loud banging sound some distance away, then my phenomenal space will change in size for as long as the banging sound continues, and shrink again once it stops. As for the intrinsic phenomenal character of the phenomenal space, it is supposed by Dainton to have this intrinsic phenomenal character in the same way as a subject's visual field might be though to have intrinsic phenomenal character, i.e. the whole of the visual field has some phenomenal character, even if they subject were looking at an all-white wall, their whole visual field would seem as of a patch of colour. [5]

I will however, not be using Dainton's notion of a phenomenal space when discussing the S-thesis. The second feature that Dainton attributes to a phenomenal space is contentious at best, and further, does not seem intimately connected to the issues at hand. What is important is that our perceptual experiences represent space in such a way that they can be spatially unified. Precisely how they do this is not the issue here. The other features that Dainton attributes to a phenomenal space, the dimensionality and egocentric axes, seem perfectly well captured by the idea of

[4] A phenomenal space is taken by Dainton to be three-dimensional in our case, i.e. everyday human perceptual experience, but may not be so necessarily.
[5] Dainton (2006)

egocentric space. For these reasons then, I will stick to talk of objective space or egocentric space when discussing spatial unity.

To return to the S-thesis as a potential explanation of phenomenal unity, if this is to succeed, it must be the case that spatial unity is both necessary and sufficient for phenomenal unity. However, as I will show through consideration of some cases developed both by Dainton and by myself, spatial unity is in fact neither necessary nor sufficient for phenomenal unity, regardless of how spatial unity itself is cashed out. Incidentally, this will put me at odds with Dainton, who considers it self-evidently true that spatial unity is *sufficient* for the phenomenal unity of perceptual experiences.

We currently have two potential versions of the S-thesis under consideration, one which concerns a spatial unity based on objective space, the other which concerns a notion of spatial unity cashed out in terms of egocentric space. For here on then, I will refer to these as the *objective S-thesis* and the *egocentric S-thesis*.

Another feature which may lend some initial plausibility to a reductive explanation of phenomenal unity in terms of spatial unity, is the S-theses' ability to explain any putative extra phenomenology that may be associated with phenomenal unity. Whilst I am still remaining neutral on the issue of whether or not phenomenal unity does involve 'extra' phenomenology, I looked in the previous section at how well a non-reductive explanation of phenomenal unity could explain any extra phenomenology were it posited, and I will do so again here, for the S-theses.

If phenomenal unity does indeed involve extra phenomenal character, and/or extra content, then the S-theses would hold that this extra must be derived from spatial unity in some way, and so the extra phenomenology would be spatial unity-specific phenomenology. One way in which spatial unity could be thought to be responsible for extra phenomenology, would be via the suggestion that there is a difference in phenomenology between spatially unified experiences and spatially disunified experiences, and that this difference is due to spatial content which is not present in the disunified experiences. This extra spatial content realises extra spatial character too, and so if phenomenal unity is thought to involve extra rather than conjoint phenomenology, the S-theorist can at least give a story as to what the extra is and where it comes from. The difference in phenomenology then, between two perceptual experiences which are phenomenally unified and two which are not, is due to spatial unity, and if there is extra phenomenology to phenomenal unity then the S-theorist is able to posit a reasonable story as to how this extra phenomenology arises.

Before I move on to consider whether or not spatial unity is necessary and sufficient for phenomenal unity, I will return briefly here to the issue of transitivity.

The non-reductive explanations of phenomenal unity which I looked at in the previous two chapters hold the thesis that the phenomenal unity relation is transitive and thus phenomenal unity cannot break down. We saw that this transitivity thesis faces some serious issues when it comes to dealing with putative counterexamples such as the split-brain cases.

It seems unlikely that such a thesis such as the transitivity thesis would be held by any S-theorist. The reason being that there seem fairly uncontroversial cases where spatial unity of one kind of another breaks down, and if phenomenal unity just is spatial unity (as it would be on a reductive explanation) then phenomenal unity can break down also. Thus, the issue of whether or not phenomenal unity is transitive or not will not be central to this chapter, as it was to the last. Instead, the crucial issue in dealing with the reductive explanations of phenomenal unity will be whether or not the kind of unity of consciousness which phenomenal unity is to be explained in terms of is both necessary and sufficient for phenomenal unity, and further, meets the nothing over and above clause.

To make a point on the satisfaction of this nothing over and above clause before proceeding, whether or not we think that an S-thesis can satisfy this clause will obviously depend on what features we take phenomenal unity to have. However, as we have seen, this is not something that there is much stated agreement on. Dainton, Bayne and Chalmers all hold that phenomenal unity is transitive, something I reject, and all received view theorists hold that phenomenal unity is a relation that holds between concrete particulars, something Tye rejects. Thus the way to find out if the nothing over and above clause can be satisfied it to look at the putative features of phenomenal unity, and see if the explanation in question can give a plausible account of them. For example, those who hold that phenomenal unity is transitive are unlikely to accept and S-thesis, as we have seen above that spatial unity is not a transitive relation.

5.5 The Objective S-thesis

If we first examine the objective S-thesis which treats spatial unity to be a matter of the objects of perceptual experience being represented as being in the same objective space,

we can demonstrate using various cases that the kind of spatial unity central to this version of the S-thesis is neither necessary nor sufficient for phenomenal unity.

To see that objective spatial unity is not necessary for phenomenal unity, we may consider cases in which perceptual experiences do not in fact have any spatial content, or at least not the right kind of content for them to be objectively spatially unified. Consider the following kind of case: suppose that at a time a subject has two auditory experiences and no other perceptual experiences. We can even suppose that the rest of her perceptual apparatus is malfunctioning in such a way that she currently has only audition. Now, suppose that as a matter of fact, auditory experiences have no spatial content. Such a claim has been made, most notably by P. F. Strawson in *Individuals*, where he claims that purely auditory experience would be experience of a 'no-space' world.[6] I will not argue for the truth of this thesis about audition here, but we can at least grant that it is possible that auditory experiences have no spatial content. An alternative, slightly weaker claim, would be that auditory experiences have no *intrinsic* spatial content, and any spatial content they do possess they do so in virtue of the other perceptual experiences that the subject has at that time having spatial content. In this way, any spatial content associated with a subject's auditory experiences is parasitic on their other experiences. If we grant at least the weaker of these two theses, then the following kind of scenario is possible: imagine a subject who has currently only two auditory experiences, which contain no spatial content, and no other perceptual experiences. We could stipulate further that in this case, all the subject's other sensory modalities are off-line. These two auditory experiences cannot be spatially unified with respect to objective space, as the subject's auditory experiences will not represent the sounds she hears as being spatially located at all. Still, despite this, it would not seem as though we have any reason to say that the subject's two auditory experiences cannot possess some conjoint phenomenology, even in this case. She may have an auditory experience as of music, and an auditory experience as of the sound of traffic, and just because there are no spatial components to her experiences does not mean that there is not something it is like for her to hear the music and the traffic together.

This kind of example makes a case for objective spatial unity being *unnecessary* for phenomenal unity. Indeed, if it works, the above example will also count against the egocentric S-thesis also, as if auditory experience is non-spatial, then it does not represent

[6] Strawson (1959). For discussion of the spatial nature of auditory experience, see Nudds (2001) and Nudds (2010) amongst others.

its objects as being located in egocentric space either. Perceptual experiences which do not have spatial content may still be phenomenally unified.

In response here, a supporter of the objective S-thesis may claim that we should not rely on examples which make controversial stipulations about perception, such as that audition has no spatial content. However, we can formulate other examples which show that perceptual experiences can lack the requisite kind of spatial content, and thus objective spatial unity is not necessary for phenomenal unity without relying on controversial claims about audition.

Suppose that a subject is currently undergoing a visual experience as of a pure white ganzfeld, an undifferentiated and uniform field of white. In this case it seems plausible that the subject's visual experience will not represent any object as being spatially located, and so the subject's visual experience will not be objectively spatially unified. However there seems no reason to say that the subject's visual experience will not be phenomenally unified.

This example does however seem to turn on the legitimacy of saying above, that a single visual experience can be phenomenally unified with itself. We have seen that Dainton's non-reductive view of phenomenal unity specifies that phenomenal unity is reflexive, and so an experience can be unified with itself. If we do not want to take this approach, that is, hold that phenomenal unity is reflexive, but instead wish to hold that phenomenal unity is not reflexive, then we could alternatively say that the subject's visual experience of the ganzfeld can be legitimately divided into more than one experience, perhaps corresponding to regions of the ganzfeld. This would still not involve any objective spatial unity, as though there may be represented spatial relations between the regions of the ganzfeld, the experiences as of different regions of the ganzfeld would still not represent their objects as being located in objective space.

The above example may be bolstered still further by supposing that as well as the visual experience(s) of the ganzfeld, the subject also has an auditory or tactile experience. Though here the object of the subject's tactile experience, say, may be represented as being located in objective space, it will not be represented as being located in the same space as the ganzfeld, due to its not being represented as being spatially located at all, and hence the two experiences will not be objectively spatially unified. Still, as above, there does not seem any reason to deny that these experiences are phenomenally unified. Here also, objective spatial unity is not necessary for phenomenal unity. An alternative case

here would be one of a subject's having an experience of darkness.[7] If it is possible to have a visual experience as of an absence, such as an experience of darkness, then this seems like a good candidate for an experience with no spatial content. Again however, there does not seem to be any reason to rule out an experience of darkness being phenomenally unified with another perceptual experience had by the same subject at that time. We should conclude that objective spatial unity is not necessary for phenomenal unity.

Can we show also that objective spatial unity is not sufficient for phenomenal unity? I think we can. Perhaps the most obvious potential example of this insufficiency would be split-brain cases.

Take a split-brain subject undergoing the standard 'key-ring' experiment performed on split-brain subjects, which was discussed in the previous chapter. The subject reports having seen only the stimulus which has been displayed on the right half of the screen ('ring') and further, denies having seen the stimulus presented to the left visual field, and denies recognising the object presented to his left hand. At the same time, the subject uses his left hand to correctly retrieve the object matching the word presented to the left visual field ('key'). When asked to name the object selected by their left hand, the will respond "ring", matching the stimulus to the right visual field. In this case the subject has a visual experience of 'key' and (though they don't report it) a visual experience of 'ring', both of which represent their objects as being located in objective space. But we have reason to think that split-brain cases can involve phenomenal disunity, and so despite objective spatial unity being present, it is not sufficient for phenomenal unity.

What reason do we have to think that there is any kind of spatial unity in split-brain cases? Well, the subject has two visual experiences, each of which represents its object (the word 'key' or 'ring') as being located in objective space (and egocentric space also). As long as the two experiences both represent their objects as being located in the same common space, then they are spatially unified. As I mentioned when introducing the notion of spatial unity, spatial unity does not require access unity, so the subject does not need to be able to jointly access the spatial content of their two experiences for them to be spatially unified. We can see more clearly then, why we can posit objective spatial unity in the split-brain cases.

[7] For more on the perception of absences, see Sorensen (2008) and Richardson (2010)

However, if we do not wish to rely on split-brain cases, the interpretation of which is controversial to say the least, as we have seen, we can present another case in which objective spatial unity is not sufficient for phenomenal unity.

Any case of a failure of transitivity could potentially be a case equivalent to the split-brain cases here. Suppose that a subject has at a time a visual experience, a tactile experience and an auditory experience. Suppose further that these three experiences are all objectively spatially unified, as they all represent their objects as being located in the same objective space. However, if it is possible to have transitivity failures in the phenomenal unity relation, then this objective spatial unity may still be accompanied by a breakdown in phenomenal unity. Objective spatial unity, if these transitivity failures are possible then, may not be sufficient for phenomenal unity.

The problem here however, is that the defender of the objective S-thesis is likely to complain along the following lines about this example: the objective S-thesis identifies phenomenal unity with objective spatial unity, and it is likely to be objected that the above example relies on transitivity failing, and if the objective S-theorist has any sympathy with the idea that the transitivity of objective spatial unity cannot fail, then this example may fail to get any purchase.

Despite potential reservations about the examples which could be used to show objective spatial unity's insufficiency for phenomenal unity, we should agree that it is unnecessary. Thus, we do not have a potential reductive explanation of phenomenal unity here, and the objective S-thesis fails.

The alternative would be to try and explain phenomenal unity in terms of egocentric spatial unity, but we can show that the egocentric S-thesis fails also, and egocentric spatial unity is also neither necessary nor sufficient for phenomenal unity. Of course, if two experiences are not objectively spatially unified, then they cannot be egocentrically spatially unified and so if the objective S-thesis falls, the egocentric S-thesis falls with it. Still, it is worth considering the egocentric S-thesis on its own also.

5.6 The Egocentric S-thesis

To see that egocentric space is neither necessary nor sufficient for phenomenal unity, we can consider the following cases. First, suppose that we could have a creature that has one body, but has two brains inside the same head. These two brains share the same perceptual apparatus, but these creatures are not like split-brain patients: there is no

integration of their separate mental states. Now it seems that the perceptual experiences that that this two-brained creature had, would be egocentrically spatially unified as they would represent objects as being located in the same egocentric space. This is because the two brains share the same body, are in the same head, and share the same sensory apparatus. Thus a visual experience as of a cup would represent the cup as being located at the same point in egocentric space, to the right of the creature's right arm, say, despite there being two brains undergoing the experience separately. Still, the experiences had by this creature would not be phenomenally unified, as there is no reason to say that there is phenomenal unity across the two separate brains, any more than there is reason to say that there is phenomenal unity between my experiences and those that are had at the same time and of the same object, by someone else. Again, it seems as thought this type of spatial unity is not sufficient for phenomenal unity.

One may respond here that this creature is not a single subject, and therefore may not have a single egocentric frame of reference, but as mentioned above, I do not think that we currently have a good enough definition of what counts as a single subject, to rule out that the creature described is not a single subject with a single egocentric frame of reference. After all, we do not generally think that split-brain patients are two subjects, with two egocentric frames of reference, or at least, we are not bound to this conclusion. This example is, as I said above, working with the traditional conception of the egocentric framework as being a torso-centred set of axes, but it seems that this version of the S-thesis will not provide for a reductive explanation of phenomenal unity, even if other conceptions of the egocentric framework were used here instead. Say we are operating with a conception of the egocentric framework on which a subject has several non-competing frameworks, centred on their limbs. On this view, we would get the same result, as we could still have egocentric spatial unity, and not have this give us phenomenal unity.

Again, it would also be possible to use split-brain cases to put pressure on a sufficiency condition made by the egocentric S-theorist. We could say of split-brain cases that they are equivalent to the above case inasmuch as the perceptual experiences of the split-brain subject may be egocentrically spatially unified without being phenomenally unified. As with objective spatial unity then, egocentric spatial unity is not sufficient for phenomenal unity.

To see that egocentric spatial unity is not necessary for phenomenal unity, we can look at an example discussed by Dainton.[8] Dainton attempts to provide examples of experiences which are phenomenally unified, despite not being spatially unified, showing that egocentric spatial unity is not necessary for phenomenal unity. Dainton's argument takes the form of a thought experiment which has its genesis in Daniel Dennett's (1979) article 'Where Am I?'. [9]

In Dainton's variation on this thought experiment, you find yourself in a scenario where your brain has been removed from your body by the kind of evil scientists who routinely prey on philosophers, and envatted in a laboratory somewhere. You are then equipped with an artificial head, complete with sensory apparatus, which is separate from your body and which can, via some sophisticated transceivers, communicate to your brain exactly the kind of perceptual experiences you would be having otherwise, were your brain in your head as normal. Now the evil scientists engineer a number of different scenarios:

1. In the first scenario, your artificial head is switched on, and placed in the audience at a loud concert, whilst your body remains in the laboratory, devoid of sensation.
2. In the second scenario, your body, which has been fitted with similar transceivers, is placed underwater and switched on, whilst your artificial head is switched off again.
3. In the final scenario, your artificial head is taken to a far-flung and snowy mountain top, and switched on, whilst your body remains switched on and underwater.

In the first of these scenarios, Dainton claims that you would be presented with a single unified space, and that your experiences would be spatially and phenomenally unified. In the second scenario, Dainton claims that again we would be presented with a single unified space, albeit this time containing both tactile and bodily experiences as opposed to visual experiences, and again we would have spatial and phenomenal unity. In the third scenario however, Dainton claims that there would be two separate spaces, thus there would be no spatial unity, though there would still be phenomenal unity. We can

[8] Dainton (2006)
[9] Dennett (1979)

say that in all three scenarios we are concerned with egocentric space, and so in this final scenario, there are two separate egocentric spaces.

So in this third scenario, Dainton suggests that there would be two spaces (it would seem most natural to think of them as egocentric spaces); one centred on your body, which is underwater, and the other centred on your artificial head, which is wobbling atop some frozen mountain. Because of these two separate egocentric spaces, there would spatial disunity in this case, but there would be no reason, so Dainton claims, to think that the subject's visual and tactile experiences were not phenomenally unified, as they would still be had simultaneously, and would still seem unified to the subject in the same way that they would were the situation more normal. Again, take the spaces in question to be egocentric spaces.

Dainton believes that the possibility of this third scenario refutes the S-thesis, and we can see how it would count against the egocentric S-thesis, showing that egocentric spatial unity is not necessary for phenomenal unity. Since it would seem to you as though you were in two places at once, Dainton claims, your audio-visual and your bodily sensations would be representing the objects of your experiences as being located in two different egocentric spaces. As these experiences would still be phenomenally unified, so the egocentric S-thesis is falsified.

The problem with Dainton's example though, is that it does not seem necessarily the case that it *would* seem to you as though you were in two places at once in this third scenario in which your body is underwater, and your artificial head has been placed on some mountain top. The alternative interpretation of what would be going on in this case would be that it may seem to you, in this scenario, as though you were in *one place*, but that you were undergoing conflicting sensory experiences, thus, there may be only one phenomenal space, which is filled with conflicting sensory information.

In response to Dainton's claims about this case, Douglas Meehan suggests that we would still have cross-modal integration of hand-eye coordination, and if we were asked to point to a bird on the mountain top, we would still point, despite it seeming to us as though our hand were underwater (this being the case).[10] Meehan claims that our kinaesthetic and visual experiences would still be 'in tune' in this way, even thought we would not end up pointing at the bird, and that this suggests that our visual experiences still represent the bird as being located somewhere relative to our hand.[11] Were our brain

[10] Meehan (2003) Note here that Meehan's replies to Dainton refer not to Dainton (2006) but to the first edition, published in 2003. All references made in this thesis however, are to (2006).
[11] Meehan (2003)

and head and body reassembled in the normal fashion, then our hand would point to the place where the bird was.

> In fact, there is further reason to think that even if you did seem to be in two places, your so-seeming would be due to something other than your bodily and audio-visual experiences' being located in separate phenomenal fields. You would, no doubt, maintain cross-modal integration, such as hand-eye coordination. When asked to point at the bird, you would point even though your finger would feel like it was underwater. Your kinaesthetic sensations would still be calibrated to your visual sensations even though you would not successfully point at the bird. This suggests your bodily sensations would be located in the same phenomenal field as your audio-visual sensations. So, if you did seem to be in two places at once, this must be due to some other factor involved in your locating yourself, not the phenomenal locations of your sensations.[12]

Meehan's claims here do carry on the talk of the sensations and experiences *themselves* as being located in the same space, but if we continue to talk only of the represented location of the objects of sensations and experiences, we can still see that Meehan's point about cross-modal integration, if true, would be damaging for Dainton, as it would demonstrate a level of perceptual integration which he wants to be absent from his example. What precisely then, would this cross-modal integration consist of? Meehan takes it that this cross-modal integration would manifest itself in an ability to point at the bird you see on the mountain-top, and at least get the direction relative to your body right, even though you would not obviously succeed in pointing at the bird. This suggests that there is still integration between the audio-visual and the bodily perceptual fields.

Dainton makes some replies to Meehan, which I will look at here, as well as points that arise from a relevant discussion between Dainton and Antti Revonsuo on the same subject.[13] This discussion serves to bring into sharper focus the key point concerning Dainton's example above - that we may have reason to suppose, contra Dainton, that in this example our perceptual fields would re-integrate spatially, rather than remaining separate, or at least Meehan argues that Dainton has not shown why the perceptual-fields would not re-integrate in this way. This is important for the current discussion, as if the proponents of either of the S-theses can show that in putative cases of spatial disunity,

[12] Ibid.
[13] Revonsuo (2003)

our perceptual fields would in fact re-integrate in a way that suggests there is no spatial disunity after all, then this would lend support to the S-theses.

5.7 Meehan & Revonsuo on Dainton

To go into more detail of the discussion between Meehan and Dainton; Meehan reconstructs Dainton's argument for spatial disunity in his thought experiment as follows:

1) In the final scenario described above, it would seem to you as though you were in two places at once: one revealed to you via your audio-visual experiences, the other revealed by your bodily experiences.
2) Since you seem to be in two distinct places, you bodily and audio-visual experiences must be located in different phenomenal spaces.
3) Nonetheless, these experiences are co-conscious
4) So, the S-thesis is false.[14]

Meehan's objections to this argument can be summarised as: A) This is not an accurate description of what it would be like for you in the final scenario, and B) Even if it was, this would not be because premise 2) is true. This Meehan takes to be so because as mentioned above, Meehan claims that in the final scenario, your audio-visual and your bodily fields would remain integrated, and because 2) supposes that our sense of location at any given time must depend entirely on our current perceptual experience, but it doesn't, as aside for locating ourselves egocentrically based on perceptual experience, we can also locate ourselves allocentrically, i.e. in Glasgow.

Dainton responds to this with the claim that both these criticisms of Meehan's are misguided, and the reconstruction of his argument is inaccurate. The actual form of his argument he states as:

1) In the final phase of the envisaged scenario, your perceptual experiences would be occurring in two distinct phenomenal spaces, one audio-visual, and one bodily.

[14] Meehan (2003)

2) Since the contents of these perceptual fields are so disparate, you would seem to have two locations at the same time.
3) The contents of your audio-visual and the contents of your bodily field are still co-conscious.
4) The S-thesis is false.[15]

It seems though, that what Dainton claims to be the actual formulation of the argument still does not escape Meehan's criticisms, as the above complaint still seems to stand against premise 2), even on the above formulation. Meehan's central point against Dainton, that there is just as much reason to think that in the final scenario our sensory fields would remain integrated as there is to think that they would come apart, still stands, as it is the 'disparate' nature of the two perceptual fields that Meehan is disputing, not the claim about your seeming to be in two places at once.

As mentioned above, what confuses matters in this discussion is both parties sticking to talk of the location of experiences themselves, and continued use of the notion of a phenomenal space. We can discard both the notion of phenomenal space, and talk of the location of experiences, and focus on the crucial point in this discussion. The crucial point here is what backing Meehan can lend to his claim that in the above kind of scenario, your experiences would not be spatially disunified, but rather would still be integrated in some way.

Meehan's reasons for making this claim seem to come from a claim he makes, mentioned above, that in the final scenario outlined in Dainton's thought experiment, where your body is underwater, and your artificial remote head is atop the mountain, you would retain cross-modal integration of your perceptual experiences. Further, this claim seems based on the claim that in this third scenario, our perceptual experiences would remain cross-modally integrated. Again however, the issue is confused here through both parties talking of the location of experiences rather than the represented location of the objects of experiences. Meehan claims that given cross-modal integration, it would not be the case that your experiences are located in different spaces. The claim that would be important for our purposes however is not this one, but a claim that cross-modal integration shows that the objects of our experiences are not represented as being located in different spaces, objective or egocentric. This claim however, Meehan does not address. He suggests that in the third scenario of Dainton's thought experiment, you

[15] Dainton (2004)

would still 'point at the bird', and so hand-eye coordination would be retained. This however, does not necessarily tell us anything about the presence of absence of spatial unity in this scenario. It may speak to the issue of whether or not we would seem to be in two places at once in this situation, but this is not what we're dealing with.

As it happens, Meehan does also reject the S-theses, but for a different reason than the ones Dainton outlines. He believes that cases of unilateral visual neglect show that spatial unity is not sufficient for phenomenal unity, as well as being unnecessary. In these cases, where patients suffering from unilateral visual neglect claim to have no conscious experience of stimuli presented to the neglected portion of their visual field, experimental data (Bertelson et al (2000)) has emerged which suggests that in actual fact, the subjects *are* influenced by stimuli to this neglected portion of their visual field. Visual stimuli presented to the neglected portion of the visual fields of the subjects apparently influences the way they respond to auditory stimuli: when asked to point in the direction of the sound, the subjects do so inaccurately, and the precise pattern of error strongly suggests they are being influenced by the simultaneously presented visual stimuli.[16]

> This explanation relies on the subjects' actually seeing the visual stimuli and on the spatial integration of visual and auditory sensations. Since the subjects see the stimuli without consciously seeing them, and since seeing them affects their
> pointing at auditory stimuli, their visual and auditory sensations must be co-present in the same phenomenal field without being co-conscious. So co-presence in a phenomenal space is insufficient for making two states co-conscious. The S-thesis is false.[17]

Dainton denies this, as he believes the sufficiency thesis to be self-evidently true. Perhaps in response to Meehan, Dainton can make use of the access/phenomenal consciousness distinction in his favour here; so that in the visual neglect case, the stimuli may be spatially access unified without being phenomenally conscious. Dainton's response is somewhat unclear, but seems to be roughly along these lines.

Despite this response to Meehan however, I have suggested cases above which also seem to show that neither objective nor egocentric spatial unity is sufficient for phenomenal unity, and these cases coupled with Meehan's claims here, put considerable pressure on Dainton's claim that the *sufficiency* of spatial unity for phenomenal unity is self-evidently true, as these examples show precisely that it is insufficient in some cases.

---

[16] Dainton (2004) p. 22
[17] Meehan (2003)

Meehan also points out a further factor which enters into the issue of self-location. He points out that our sense of location is not simply determined by our current perceptual experiences, but also be more general beliefs about our allocentric location, that I am in Glasgow, for example. Current perceptual experience does not automatically trump these background beliefs. Meehan suggests that in the final disputed scenario in Dainton's example, if you do seem to be in two locations at once, it is because of these general background beliefs and not due to your current perceptual experience.

Dainton responds that there is something amiss in Meehan's remarks about self-location: as we have seen, in the disputed third scenario, Meehan suggests that you may seem to be in one place, with conflicting sensory information, and if you did seem to be in two places at once, this would be due to something other than the state of your current perceptual experiences. Dainton takes issue with both of these claims, though it is to the first that his response is the most incredulous: Dainton can see no reason why you would take either your audio-visual, or your bodily perceptual experiences to be hallucinatory, thus it seeming to you that you were simply in one place, but with conflicting sensory information. Dainton concedes that you might come to the conclusion that you were standing on the mountain, but undergoing hallucinatory bodily experiences as of being underwater, if you weren't aware that the experiment was going on, but you are, claims Dainton, so you would not make this assumption of hallucinatory or illusory experiences.

> As a consequence, you have no reason to believe either of the sensory channels is hallucinatory. The causal chain from physical stimulus to perceptual experience is certainly of a non-standard kind, but — and as you well know — it is of a kind which is providing you with reliable information about your environment(s). What we make of our experience at a give time depends — as Meehan himself notes — on our more general beliefs about our predicament. The case in question provides an excellent illustration of this very point.[18]

It is not obvious from what Dainton says, however, that your background beliefs in this case *would* change your experiences: despite having the background belief that you were undergoing this experiment, it may still *seem to you* as though you were on top of the mountain, but with the bodily experiences as of being underwater, even if, due to your background beliefs about what is going on, that you believe yourself to be back in the

[18] Dainton (2004) p. 23

laboratory.[19] This issue of how your background beliefs would affect how it would seem to you in a scenario like this will arise again later. Dainton also claims that if you do seem to be "bi-located" in the third scenario, it would be due to a combination of both factors: current perceptual experiences, and background beliefs about your location: it is because of your background beliefs about the experiment that you would not regard your current experiences as illusory, and it is because of these experiences that you would seem to be bi-located. Still, on the basis of your perceptual experiences alone in this case, how it would seem to you may go either way.

The crux of Dainton's reply seems to be the claim that Meehan is wrong to think that Dainton's interpretation of the third scenario depends on it seeming to you that you are in two places at once, when Dainton claims the crucial fact is the spatial isolation of your audio-visual and bodily sense-fields. This seems to trade on the more controversial/confusing reading of the S-thesis, where it is the experiences themselves that are located in the same phenomenal space, rather than the more coherent notion that it is the objects of those experiences. If it is the former, then in virtue of the fact that body and head are separated, then surely it will be trivially true that the bodily and audio-visual experiences are dislocated. Again, this arises from thinking of the S-thesis as involving claims about where the experiences are located, as opposed to what I have claimed is the more natural claim about where the *objects of experience* are represented as being located.

However, in elucidating the above claim, Dainton claims that there are three ways of answering the question 'where are you?':

> *Subjective location*: where you seem to be *qua* thinking subject; where your conscious thinking, remembering and imagining seem to be occurring.
> *Sensory perspective*: the location from which you are perceiving the world.
> *Material location*: the location of the material basis of your consciousness, i.e. your brain or nervous system.[20]

Under normal circumstances, Dainton claims, all of these will coincide, but they can come apart, as in the experiment Dainton outlines.

---

[19] This case may be an example of background beliefs failing to alter illusory perceptual experiences; just as the Muller-Lyre illusion persists even though we believe that the two lines are the same length.
[20] Dainton (2004) p. 24

I feel however, that this debate misses the important point somewhat: what's at issue is whether or not my experiences are phenomenally unified *in virtue of* being spatially unified, not whether a certain conception of spatial unity could result in our seeming to be in two places at once.

Antti Revonsuo is also unconvinced by Dainton's description of what is going on in the disputed third scenario.[21] Revonsuo also cites putative empirical findings in order to put pressure on Dainton's conception of what happens in the third scenario, where your artificial head is on top of the mountain, and your body is underwater. He claims that findings from experiments on binocular fusion/rivalry show Dainton cannot be right about what would be going on in his third scenario. When two different images are shown to a subject's eyes at the same time, then either the images are merged together, if they are sufficiently similar (binocular fusion), or only one of the images will be perceived at any one time, if they are not (binocular rivalry). Revonsuo believes that this effect generalises to consciousness as a whole, and so claims that either fusion or rivalry would be evident in the third scenario.

Dainton disagrees. He claims that there are two general problems with Revonsuo's claims. The first is the extrapolation from the purely visual cases of binocular fusion and rivalry, to consciousness as a whole. He claims that while these claims made about the visual system sound plausible, there is little evidence to suggest that an analogous effect happens *across* sensory modalities, as in the third scenario, and that there would be problems in reconciling conflicting content from the visual and the bodily sensory modalities. Dainton says:

> If I listen to the sound of heavy rainfall while wandering about on a dry sunny day — I am using a personal stereo — my brain makes no attempt to impose thematic harmony on my audio-visual experiences. I may *hear* rain falling, but I do not start to *see* rain falling — my visual experience remains largely unchanged in response to dramatic alterations in what I am hearing. Not only is there little or no fusing, there is no alternation and suppression: I continue to have both auditory and visual experiences continuously. A similar independence exists between audio-visual and bodily experience. If while walking though a hot desert I were to don a virtual reality headset which provides me the kinds of sights and sounds I might expect to have while swimming underwater, I would not

[21] Revonsuo (2003)

> suddenly feel as though my body had been plunged into water — I would continue to feel the hot dry air against my skin, not cold, heavy, water.[22]

So Dainton claims, but is this the case? Surely the McGurk effect and other cross-modal effects demonstrate that there is some cross-modal influence, and that changes in auditory and visual experience can affect each other.[23] Dainton does talk about the ventriloquism effect, where a sound is experienced as coming from an object that is not the actual sound source, but seems to be suggesting that this is too limited an effect to trouble him, though he says little to back up this claim.

The second of Dainton's problems with Revonsuo's claims revolves around the second claim, that if in the third scenario, you are undergoing both audio-visual and bodily experiences, then they would somehow be spatially superimposed. Dainton concedes that this may be possible, but claims it would be no more likely than spatial disunity in this case. Again, after the Dainton/Meehan discussion, it seems that Dainton concedes that this explanation is just as likely as his own interpretation as to what would happen in this scenario.

It seems then, as though Dainton's example provides a somewhat less than conclusive refutation of the S-thesis. I think however, that examples of the kind that Dainton wants here are possible to provide, and so I will introduce my own example which successfully serves to refute the S-thesis: the Big-Numb-Arm.

5.8 The Big-Numb-Arm

Imagine that one of your arms is incredibly long and flexible, and that it is missing all senses of proprioception and kinaesthesia, as well as tactile sensation in the arm. You retain tactile sensation in your hand however and as such retain the ability to identify objects via exploratory touch. It seems as though you could become confused as to where your hand was located, relative to your body, or to the objects in your perceptual field. It seems also as if you could not point, with the hand of the big-numb-arm, to an object that you can see or the source of a sound that you can hear. So, imagine now that

[22] Dainton (2004)

[23] The McGurk effect is thought to demonstrate a cross-modal interaction between vision and audition. Subjects are initially shown a video of a person pronouncing a phoneme, such as 'ba' and played an audio track which syncs up with the video, but are then played an audio recording of someone pronouncing a different phoneme from that which is being pronounced in the video. In this case, subjects tend to hear a third, intermediate sound; a visual 'ga' combined with an auditory 'ba' is usually heard as 'da'. (McGurk and MacDonald (1976))

you are currently having a visual experience as of various objects on a table in front of you, and simultaneously, a tactile experience as of holding a cup in the hand of your big-numb-arm, the location of which you are unaware of. Here, it seems as though we would have a genuine case of two separate phenomenal spaces, and as there seems no reason to think that there would be phenomenal disunity in this case, we have phenomenal unity without spatial unity. In this big-numb-arm case, we have one phenomenal space, which corresponds to your auditory, visual, and various other experiences, and another, separate phenomenal space which corresponds to the tactile experiences you have via the hand of your big-numb-arm. We have then, a case of phenomenal unity without spatial unity, which shows that spatial unity is not necessary for phenomenal unity.

There are several points which count against Dainton's putative examples of phenomenal unity without spatial unity which do not affect the Big-Numb-Arm, and I believe also that the objections of Meehan and Revonsuo will not have as much bite with regards to my example as they have with regards to Dainton's.

In defending the Big-Numb-Arm as a better example, a crucial point, mentioned in the description of the case, is that it would not be the case that you could make any indication as to where the hand on the end of your numb arm was located, relative to you or the rest of the objects that you were perceptually aware of at the time. This avoids Meehan's objection to Dainton, where Meehan claimed that in Dainton's third scenario, you would still be able to point at where you took the bird to be relative to your body, despite your body being underwater. In the Big-Numb-Arm case, this seems impossible, as you have no idea where the hand of your numb arm is located.

Another point where the Big-Numb-Arm example avoids the problems of Dainton's case, is that the objections raised by Meehan and Revonsuo, concerning the possible perceptual field integration that would go in Dainton's third scenario, do not seem to affect the Big-Numb-Arm. Much of the debate between Meehan and Revonsuo on the one hand, and Dainton on the other, was over whether or not, in the third scenario of Dainton's thought experiment, your bodily, and your audio-visual sensory fields would exhibit the kind of dissociation that Dainton suggests. The Big-Numb-Arm example more successfully avoids this issue altogether, as in this case, your bodily and your audio-visual sensory fields *would* be integrated, all but the tactile experience you have via the hand of your numb arm, which would remain spatially disunified from your other tactile experiences, and your other audio-visual experiences. I think that this means that the Meehan/Revonsuo objections have less bite here. Nor do I think that it matters that

the spatial disunity may be thought to be confined to tactile experiences, as even if this is the case, it would still count as an example of phenomenal unity without spatial unity.

Another objection that applies to the Big-Numb-Arm case, and may seem to have more bite, goes as follows: suppose that I have left the hand of my numb arm in my office, where it is grasping a very familiar object; say a paperweight, that is on my desk. It may be thought that through recognising the paperweight, I can come to know where my hand is located, and that any spatial disunity in my experiences may evaporate. I do not think however, that this would happen in this case. Despite coming to have a belief about where my hand was located, this would not affect my experiences of the hand being spatially disunified from my other perceptual experiences. It would not seem to me that the paperweight was occupying the same phenomenal space as the other objects of my experiences, and so spatial disunity continues, even with background beliefs about where my hand is. So, even though I may be able to locate my hand allocentrically, I would not be able to locate it egocentrically, or to put it another way, I may be able to locate my hand in objective space, but it would still not seem to be in the same egocentric space as the other objects of my experience. I think that this goes some way to showing why the Big-Numb-Arm would always give rise to two disconnected egocentric spaces, the frame of reference of one of which is centred on, and restricted to the movements of the one hand.

Can there be any kind of response to the big-numb-arm on behalf of the egocentric S-theorist? It could be suggested that the object represented by touch via your big-numb-arm hand, is represented as bearing an indeterminate spatial relation to your body, and the objects of your visual experiences, even though the exact location is not represented. This would yield a weakened version of the S-thesis, which would run something like as follows: simultaneous experiences are phenomenally unified in virtue of their objects being represented as being located in the same space, despite not necessarily being represented as being located in the same egocentric space, and despite the exact spatial relations not being represented. So, if your big-numb-arm was half a mile long, then you might represent the cup as being an indeterminate distance and direction from your body, but still located in space within a half-mile radius of your body. The problem with this response, and the weakened S-thesis, is that it seems to leave an S-thesis that is *too* weak. We are left with something that seems close to being vacuous, and without a determinate representation of spatial locations. This weakened S-thesis would perhaps also fall foul of the same problems as the thesis based on the objective space reading of

spatial unity, as without determinate spatial relations being represented, we are left with the objects being located merely in objective space somewhere, and we have seen that this kind of spatial unity is neither necessary nor sufficient for phenomenal unity.

Before making my conclusions, I will make a point about the status of the thought experiments which have been used in this chapter. I am using the big-numb-arm and the other thought experiment-based cases to make a *conceptual point*; that phenomenal unity and spatial unity can come apart. If empirical research shows that analogous cases to the big-numb-arm or the other thought experiments I have outlined in this chapter could actually happen, then so much the better, but I am also happy if it turns out that something like the big-numb-arm case is not nomologically possible, as this will not affect the conceptual point that I am making. The reason for this is that any proponent of the S-thesis, on either of the variations that I have looked at in this chapter, is using it to try and reduce phenomenal unity to spatial unity of some kind. In doing this, the S-theorist is making an identity claim, they are claiming that we can reduce phenomenal unity to spatial unity, and the two are thus identical.

So, in order to refute an identity claim such as this, showing that phenomenal unity and spatial unity are in fact conceptually distinct is sufficient. The cases in which phenomenal unity and spatial unity putatively come apart do not have to be nomologically possible, that is, possible according to the laws of nature on this world, but need only be possible on some world.

5.9 Conclusion

So to conclude, we have seen that despite initially seeming to be a promising avenue for enquiry, vis a vis a reductive explanation of the phenomenal unity of perceptual experiences, the attempt to reduce phenomenal unity to spatial unity fails. We have looked at two variants on this S-thesis, one that is cashed out in terms of objective spatial unity, and one that is cashed out in terms of egocentric spatial unity. I presented various cases which I claimed should lead us to conclude that neither of these kinds of spatial unity is either necessary or sufficient for phenomenal unity: objective spatial unity is not necessary for spatial unity, as we saw from cases of phenomenally unified experiences, one or more of which did not have the kind of spatial content necessary for the subject's experiences to be phenomenally unified. Objective spatial unity was shown to be insufficient for phenomenal unity through the split-brain cases, where the split-brain

patient's experiences may be objectively spatially unified, despite not being phenomenally unified.

Egocentric spatial unity was also shown to be neither necessary nor sufficient for phenomenal unity. Again, split-brain cases are potentially cases where the subject's experiences are potentially egocentrically spatially unified, without being phenomenally unified. We also saw the same thing would be the case in the hypothetical case of the creature with two brains and a single body. The big-numb-arm case demonstrates egocentric spatial unity's being unnecessary for phenomenal unity.

From all of this, we are in a position to reject this potential reductive explanation of phenomenal unity, and conclude that phenomenal unity cannot be successfully reduced to spatial unity. In the next chapter I will move on to look at the possibility of explaining phenomenal unity in terms of introspective unity.

Chapter Six - Phenomenal Unity as Introspective Unity

6.1 Introduction

We have seen in the previous chapter that a reductive explanation of phenomenal unity in terms of spatial unity does not succeed, despite the initial plausibility of such a thesis. In this chapter I will consider an attempt to reduce phenomenal unity to another kind of unity exhibited by consciousness, introspective unity. This explanation has been quickly dismissed by both Tye and Dainton as we shall see, but I demonstrate why the versions of this reductive explanation posed by Tye and Dainton are unsatisfactory, and will construct and examine more plausible versions. Having said this I will conclude, after looking at these alternative introspective unity theses, that ultimately this also fails as a reductive explanation of phenomenal unity and the reason for this failure is the commitment to it being introspection in some form that does the explanatory work. In short, although there are better versions of the introspective unity theses than the ones Tye and Dainton provide, there is no successful reductive explanation of phenomenal unity to be had here.

This chapter will go as follows then: firstly, I will provide some brief background discussion of the various theories of what introspection is, and discuss how the differences between them would affect a potential explanation of phenomenal unity in terms of introspective unity. I will then move on to discuss the idea that phenomenal

unity can be fully reductively explained by introspective unity, or joint introspectibility. As with the reductive explanation looked at in the previous chapter, the idea of explaining phenomenal unity in terms of introspective unity has some initial plausibility, which I will highlight. Next I will outline the various forms this explanation in terms of introspection could take. This reductive explanation in terms of introspection is quickly rejected by Tye, Dainton and by Cody Gilmore and I will consider the arguments they make in support of their rejection.[24] The approach they use is similar to that used to reject the spatial unity theses considered in the last chapter: they seek to demonstrate, via consideration of various cases, that introspective unity is neither necessary nor sufficient for phenomenal unity.

In contrast to the previous chapter however, the arguments against the introspective unity explanation of phenomenal unity are not so convincing, directed as they are at an implausible version of the thesis, and something of a straw man. However, despite it being possible to explain away some of the initial objections to the I-theses, I will show that the commitment to introspection in some form being what explains phenomenal unity will always compromise any version of the I-thesis, and so it should not be endorsed. Before all this however, I will give a very brief overview on the discussion of introspection itself, and what sort of thing we should think it is.

Before providing some background to the various theories of introspection, I will outline again here what is meant by introspective unity, or joint introspectibility. As mentioned in the introductory chapter, and above, introspective unity concerns the joint introspectibility of two or more conscious states (here again I am going to limit my discussion to perceptual experiences). These states will be introspectively unified if and only if they (or their contents) can be introspected together simultaneously in a single act of introspection. This single act of introspection will involve the two states being simultaneously introspected by the same subject and further will need to involve some kind of higher-order state which contains the content and character of both the introspected states. Of course what this 'single act of introspection' is will also depend in part on what you think introspection itself is; whether it is some kind of 'inner sense' akin to perception, or some kind of causal process. I will discuss this further after introducing the main theories of introspection, and will suggest what a single act of introspection should be taken to mean on each of the main accounts. Another key point here, which will be a pivotal issue in this whole chapter, is the difference between *actual*

[24] Tye (2003), Dainton (2006) and Gilmore (2003)

and *potential* objects of introspection, i.e. the difference between a state's being *introspected*, as opposed to *introspectible*. This difference will also affect the potential explanation of phenomenal unity in terms of introspective unity, as if for two states to be introspectively unified, they have to *actually be introspected* together simultaneously, then this condition seems fairly obviously unnecessary for phenomenal unity. This is one of the failings with the thesis that Tye and others criticise, but as I will claim, there is no need to formulate the claim in this way, and there would be little motivation for doing so.

As an important note here, it is important to establish what is meant by 'higher-order' in what will follow. On the various theories of introspection I will shortly outline, the introspective state is usually described as being a 'higher-order' state. In this context, 'higher-order' is to be understood as denoting some kind of iterated intentional operator. Take the example 'John wonders whether he knows that Goldbach's conjecture is true.' Here the wondering is a higher-order state. 'John believes that he is having a perceptual experience as of a bird' involves a higher-order introspective state, the belief that he is having a perceptual experience. The belief here is higher-order insofar as it is 'about' the perceptual experience. The same will be the case even on non belief-based theories of introspection: the introspective state will be 'about' the introspected state, and so will in this sense be a higher-order state.

6.2 Introduction to Introspection

There are various competing theories of what kind of thing introspection is, with one big debate here being over whether or not introspection is analogous to perceptual experience in any way. Though I will not take a stance on which is the correct theory of introspection, it will be useful to see precisely what they are, especially as later in this chapter, different introspective unity explanations of phenomenal unity will be described which invoke different views on introspection to various ends. In this section I will highlight the debate between causal and non-causal accounts of introspection, part of which is the debate between so-called 'inner-sense' views of introspection and their opponents.

An account of introspection has to account for two features that are standardly taken to be indicative of the kind of access we have to our own mental states. These two features are *peculiar access* and *privileged access*. How a theory of introspection accounts for these two features varies according to what kind of account it is.

Once again, here I will be dealing with the introspection of perceptual experiences only (though I may still use the phrase 'mental states'), and so at points where some would claim there are differences between the introspection of perceptual experiences and the introspection of beliefs and other cognitive states, or that there is in fact no introspective access to such cognitive states, I will skip over these issues, and limit myself to talk of introspecting perceptual experiences only.[25]

To return to the kind of access introspection is taken to give us to our mental states, introspective access to our mental states is *privileged*, as we are far more reliable in gaining knowledge of our own mental states than we are at gaining knowledge of the mental states of others. Alex Byrne describes this privileged access as follows:

> [B]eliefs about one's mental states acquired through the usual route are more likely to amount to knowledge than beliefs about others' mental states (and, more generally, beliefs about one's environment). At any rate, knowledge of one's own mental state is more likely when the state is neither factive nor object-entailing. One may well falsely believe that the cat is indoors; hence one may well falsely believe that one *knows* that the cat is indoors or *sees* that the cat is indoors. Similarly, one may well falsely believe that one sees *the cat*. But it is harder to err in believing that one *believes* that the cat is indoors, or that it *looks to one* that the cat is indoors.
>
> To say that we have privileged access is not to say that beliefs about one's present mental states always amount to knowledge. Such beliefs need not even be true. One can falsely believe that one is angry, that one wants a beer, that one believes that one is happy, for example. More controversially, one can even falsely believe that it looks to one that something is red, or that one has a headache. Nonetheless, although error may always be a possibility, in a typical situation it is easier to be right about one's (non-factive, non-object-entailing) mental states (that one believes that the cat is indoors, say) than about the mental states of another (that Fred believes that the cat is indoors), or the corresponding tract of one's environment (that the cat is indoors).[26]

So, introspective access is privileged inasmuch as it forms a more reliable process than those on the basis of which we form beliefs about things other than our own mental states. Peculiar access on the other hand, the second feature that introspection is taken to have, grants us access to our mental states in a way that we do not have with the mental

[25] Peter Carruthers in his (2010) argues that there is no such thing as introspection of judgements or decisions, and rejects the introspection of "propositional attitude events" more generally.
[26] Byrne (2005) p. 2

states of others. There is some particular method by which we have access to our own mental states, unique to the process of introspection. It is this the precise nature of this method that is the source of much of the debate here.

The different ways of accounting for these two features are what marks out different theories of introspection. I will start this overview of the theories of introspection by looking at causal accounts of introspection, including those which take introspection to operate via some 'inner sense'. Much of this ground concerning the inner sense views has been extensively covered, and in many cases has been left muddied, so here I will stick to the discussions of the inner sense view presented by Alex Byrne and Brie Gertler.[27] As Byrne points out when introducing the notion, the idea that we have some kind of perception-like apparatus that scans our mental states, illuminating their contents does have a certain degree of intuitive appeal, at least when left at a certain level of generality and abstraction. However, of the kinds of metaphors commonly associated with the inner sense theory, those of the 'inner eye', should not be taken very seriously, as there are in fact no introspective experiences, nor is there any organ of introspection, as there are other perceptual organs. What is a better description of the inner sense view is the idea that introspection is the operation of some kind of internal scanning mechanism which functions to monitor our mental states. The inner sense view is also a *causal* account of introspection, in that the relation between the introspecting and the introspected state is held to be a causal one: the state being introspected plays a causal role in the introspecting state's coming to have the contents (and I use this word quite generally here) that it does have. So, as a general characterisation of the inner sense view, introspection involves the internal scanning of our mental states via some kind of causal process whereby the contents of the introspecting, or higher-order state are causally determined by the contents of the states being introspected. Described like this, the inner sense view seems less reliant on some obscure inward-looking faculty, and closer to other causal accounts of introspection. Indeed, as Brie Gertler points out, you may think that this is the only real point of analogy between introspection and perception, that they both involve some causal process. This, as Gertler says, is also the main advantage to a causal account of introspection, on which it has some structural similarity to perception: just as we can explain the relations that hold between objects in the world and states of our perceptual systems, we can explain the relations that hold between the states of our perceptual systems, and our higher-order introspective states

[27] Byrne (2005) and Gertler (2009)

by the appropriate causal relation. This provides the best approach if we are to have a naturalistic approach to introspection: it is a causal process, similar in some ways to the causal processes involved in our perceptual experiences.[28]

Much of the general objection to the inner sense view comes from the view that the dissimilarities between perceptual processes and introspection are too great, and that further, a perception-like, *causal* model of introspection does not give us what we want of a theory of introspection.

Under this further complaint is the idea that introspection seems necessarily first-personal (see the comment about privileged access above), and the worry that causal theories cannot provide this. It seems inconceivable that I could have introspective access to someone else's mental states, or that someone else could have introspective access to mine. However, causal views, and by extension inner sense views, may not satisfy this complaint. Brie Gertler outlines this objection as follows: suppose that in the future we have managed to identify the neurophysiological states which perfectly correlate with a whole range of thoughts, including all thoughts about ice cream. Now, the ubiquitous 'Evil Scientist' could now hook up some kind of monitoring device, which unbeknownst to me, relayed to him any ice cream thoughts which I had, and further, *caused him* to think 'Stuart is thinking about ice cream', when and only when I was in the relevant state. This process would be very similar to introspection, on a causal account: it would be a causal and non-inferential process, and could conceivably be as reliable as introspection.

> There are, of course, differences between the agent's awareness of my 'ice cream' thoughts, and my own introspective awareness of them. For instance, only the former depends on an artificial device. But these differences seem too insubstantial to undergird the seemingly profound difference between introspective knowledge and other-knowledge.[29]

This objection shows that causal views face a problem when it comes to keeping the privileged access to our own mental states that introspection is taken to provide us with. Of course it is open to the causal theorist to drop this demand for privileged access, and to concede that the kind of first/third person asymmetry that introspection is usually taken to have, can be dropped also. According to Gertler however, most defenders of

[28] Gertler (2009) p. 388
[29] Gertler (2009) p. 387

causal accounts of introspection would not make this concession, and would maintain the stance that there could not in principle be cases where a third person could introspect your mental states. One way in which causal theorists defend this, is by appeal to the presence of indexicals in introspective reference.

> In ordinary use, indexicals such as 'I' and 'here' refer to oneself and one's location; without further qualification, they cannot refer to another person or place. The idea, then, is that introspection involves a type of indexical reference that is similarly tied to oneself (or one's mental state). The naturalistic benefits of causal accounts will be preserved if this simple semantic feature can fully explain why one can introspect only one's own states.[30]

I shall not pursue this line here, but will instead move on to give a brief overview of non-causal accounts of introspection. Non-causal accounts operate on the assumption that a causal process can always go awry without our noticing: various things can interfere with the causal process between object in the world and our visual system, unseen filters, visual illusions of various kinds, etc. This is attributed to the 'logical space' between cause and effect. If there were no such logical space between our first-order mental states, and our introspective states, then there would be no chance of introspection going awry, or give false information about the states being introspected. Gertler identifies two distinct types of non-causal account of introspection, which are not competing accounts, but focus on the introspection of different types of state. Since the first type of account focuses on thoughts, and cognitive states, and since my focus here is on perceptual states, I will go over it only briefly. This first position has been advanced by Tyler Burge, and involves the introspecting state *incorporating* the thought 'that P'. The introspected state then is somehow incorporated into the higher-order introspecting state.[31] The second non-causal account also makes use of this incorporation relation, but applied to perceptual experiences and sensations. This view is held by Gertler herself, as well as David Chalmers.[32]

The main focus of this chapter however, is not to elucidate all possible theories of introspection, but to examine the possibility of explaining phenomenal unity in terms of introspective unity, or joint introspectibility. The purpose of the previous section has

[30] Gertler (2009) p. 387
[31] Burge (1988)
[32] Chalmers (2002) Gertler (2001).

been to give a better idea of what the mechanisms of introspection are taken to be by the various parties working on introspection. The issue of how exactly we should think of introspection will re-emerge in the next chapter also. In this next section I will also mention how adopting different views of introspection may affect the potential reductive explanation of phenomenal unity under discussion.

6.3 Phenomenal Unity as Introspective Unity

Recalling how it was described above, two states are introspectively unified when and only when the two states can be introspected together simultaneously in a single act of introspection. Here I will look at how introspective unity may be used to explain phenomenal unity, and the initial plausibility of such an explanation.

The difference in what introspective unity involves then, depending on your theory of introspection, will be that a single act of introspection will either involve a single operation of the introspection mechanism, and will mean that two states simultaneously introspected will be jointly related to an introspective judgment via the appropriate causal relation, or the two states simultaneously introspected will both be embedded in the same introspective state at the same time by the logical relation of incorporation that Gertler and Chalmers posit. As for the substantive difference between these two options, this seems to boil down to the different relations the theories posit how the introspecting is taking place: on one hand there is a causal relation, as explained by the causal accounts, and on the other, there is the more mysterious incorporation relation which features in the accounts of Gertler and Chalmers. Further potential difference between the various accounts of introspection will concern the kind of introspective state that is formed. At this stage, I will not advocate one theory of introspection over another, and as I don't think that much practical difference is made to the discussion of the potential reductive explanation of phenomenal unity by the differences between the theories of introspection, I am not going to dwell on this issue here, though I will return to it towards the end of the chapter. I will then proceed to move on to look at the idea that introspective unity is both necessary and sufficient for phenomenal unity.

Can phenomenal unity be explained solely in terms of introspective unity? I will go into more detail as to exactly what this thesis will look like shortly, but how plausible an

idea does it look initially? We saw in the last chapter that a reductive explanation of phenomenal unity in terms of spatial unity has a good deal of intuitive plausibility. Does an explanation of phenomenal unity in terms of introspective unity offer the same? Again, there seem on first sight good reasons for thinking that phenomenal unity and joint introspectibility are intimately connected. Take once more the 'key-ring' experiment performed on the split-brain subjects. Perhaps, if we are tempted by the thesis that phenomenal unity breaks down in these cases, a natural way of explaining why phenomenal unity has broken down is to say that it has done so because the subject's two visual experiences are not introspectively unified under these experimental conditions at least. The subject cannot jointly introspect both 'key' and 'ring' and perhaps this is why there is a breakdown in phenomenal unity here.

So, an explanation of phenomenal unity in terms of introspective unity can at first blush give an explanation of some problem cases. Is there any other reason why we initially take this as a plausible reductive explanation of phenomenal unity? Well, we might think that since phenomenal unity involves conjoint (or extra) phenomenology, what is responsible for this conjoint phenomenology is that the experiences had by the subject at that time were all jointly introspectible. From here on, I will refer to the claim that we can explain phenomenal unity solely in terms of introspective unity as the *I-thesis*.

To make two final points in this section, the result of explaining phenomenal unity in terms of introspective unity will have the result that phenomenal unity is not a transitive relation, and can break down within a single subject. This is because introspective unity can break down, a point which should be uncontroversial. As with the S-thesis in the previous chapter then, it will not be a feature of the I-thesis that it gives us a transitive phenomenal unity relation. Again, for those who would hold that phenomenal unity is a transitive relation, any I-thesis is going to struggle to satisfy a nothing over and above clause, much as we say with the S-theses.

Finally, the I-thesis seems reasonably well-suited to explaining any extra phenomenology that may be claimed to come along with phenomenal unity. If we were to hold that phenomenal unity involves extra phenomenal character, then this extra character would come from the experiences in question being jointly introspectible. Though this may not be quite as clear a mechanism for imparting extra phenomenology as would be there in the reductive explanation in terms of spatial unity, the I-thesis is still in a better position to explain any extra phenomenology than the non-reductive views. Once again here, I am not taking a stance on whether or not conjoint phenomenology

does involve extra phenomenology, merely noting how well the various explanations of phenomenal unity would explain the positing of extra phenomenology.

6.4 The I-thesis

A version of the I-thesis is formulated by Dainton, who discusses and rejects it in similar fashion to his treatment of the S-thesis which I looked at in the previous chapter. Dainton rejects the I-thesis, as do Michael Tye and Cody Gilmore, and I will look at their reasons for doing so in due course. First, we need a formulation of the I-thesis. Dainton provides two variations, which he calls the strong and weak I-theses. These are as follows:

> *Strong I-thesis* – Co-consciousness is constituted by introspectibility: experiences are co-conscious because they are introspected or introspectible. A group of token experiences are co-conscious if and only if they are either the actual or potential objects of a single introspective awareness.[33]

And the weak I-thesis:

> *Weak I-thesis* – Co-consciousness is not constituted by introspectibility, but the two are correlated: if a group of experiences are co-conscious they are all actual or potential objects of a single introspective awareness.[34]

In what follows I will question these formulations, but they will serve their purpose for the moment. The I-thesis then, states that two or more perceptual experiences are phenomenally unified if and only if they are the actual or potential objects of a single act of introspection. As we will see, much of the problem with the criticisms of the I-thesis stem from a failure to keep 'actual' and 'potential' suitably separate and thus to adequately distinguish between jointly introspected and jointly introspectible.

6.5 Against the I-thesis: Tye, Dainton and Gilmore

[33] Dainton (2006) p. 35
[34] Dainton (2006) p. 35

Michael Tye is quick to dismiss the I-thesis. He claims that phenomenal unity can exist independently of introspective attention, just as the experiences themselves can.[35] Introspection does not create phenomenal unity, claims Tye, it merely reveals or discloses phenomenal unity. Thus phenomenal unity and introspective unity cannot be identified. Tye's quick dismissal of the I-thesis highlights what will be problem with the other arguments against the I-thesis that I will look at here: not enough care is taken to keep apart the notions of being *introspected*, and being *introspectible*. Though Tye is right to say that the phenomenal unity of a set of perceptual experiences is independent of those experiences actually being introspected, he does not address the claim that phenomenal unity is a matter of experiences being jointly *introspectible*.

That phenomenal unity can be explained by introspective unity is rejected by Dainton also. Dainton claims that introspective unity, like spatial unity, is neither necessary nor sufficient for phenomenal unity. However, as I have mentioned above, I will show that the version of the I-thesis which Dainton outlines and rejects here is only one possible way in which introspection could explain phenomenal unity, is contingent on a kind of attention-dependence of the phenomenal character of our perceptual experiences holding, which we have reason to doubt, and overall is not the strongest version of the I-thesis that can be formulated. So although Dainton may successfully refute the version of the I-thesis that he presents, this is not the best possible version of the I-thesis, and is not one which would meet with much support even from a defender of another I-thesis. I will show that the same is true of Cody Gilmore's treatment of the I-thesis also, and so at the close of this chapter, it will emerge that a better formulated version of the I-thesis is still possible, and it will be one that is not affected by the criticism that I will look at in this section.

As outlined above, Dainton describes the I-thesis as coming in a strong and a weak variant. These two versions of the thesis are outlined above.

The immediate question raised by Dainton's formulation of two I-theses however, is what the difference between them is. The strong version of the thesis states that phenomenal unity is *constituted* by introspective unity, and that two experiences are phenomenally unified *because* they are jointly introspected or introspectible. On the strong I-thesis, all there is to two experiences being phenomenally unified, is their being introspectively unified. This sounds like the kind of thesis you would expect from a putatively reductive account of phenomenal unity. The weak thesis on the other hand,

[35] Tye (2003) p. 19

states that phenomenal unity and introspective unity are correlated, rather than the former being constituted by the latter. So it would seem that this weaker version of the I-thesis does not provide the kind of reductive explanation of phenomenal unity that the strong version does: the strong version explains phenomenal unity fully, in terms of introspective unity. When two states are jointly introspected/introspectible, that is all there is to phenomenal unity. The weak thesis however, only provides a guide by which we may pick out phenomenally unified states: they are correlated with the states which are jointly introspectible. As it stands, it is still not clear what the difference between these two theses is: if the correlation in the weak I-thesis is taken as a necessary connection, then phenomenal unity is necessarily correlated with joint introspectibility, and thus we would not have any phenomenally unified states which were not jointly introspectible. This may then be equivalent to the strong version of the I-thesis. Further, it would be possible to claim that this seeming correlation is in fact identity, thus collapsing the weak version of the I-thesis into the strong version. Of course, this last move may be denied, but then we would be left with a thesis which would tell us only about phenomenal unity's being correlated with joint introspectibility, which, although it would have epistemological value and would provide evidence for phenomenal unity, would not serve as a reductive explanation of phenomenal unity, which is the topic of this chapter. Thus, I think that we should from now on speak of Dainton as discussing one I-thesis, which is equivalent to the strong I-thesis that Dainton explicates above, and states that *two experiences of a single subject are phenomenally unified if and only they are jointly introspectible, and where this joint introspectibility exhausts phenomenal unity.*

Dainton rejects both the strong and the weak versions of the I-thesis which he outlines, though given what I have said above about the difference between the two, I shall continue to treat them as one thesis, and evaluate Dainton's arguments as they would apply to that one thesis.

Dainton's initial dismissal of the I-thesis revolves, as Tye's does, around the suggestion that phenomenal unity is independent of introspection, and though introspection can reveal the phenomenal unity of our perceptual experiences, it does not constitute it.

> The unity of the phenomenal background [which is the sum of our currently un-introspected experiences according to Dainton] is something we are passively aware of; although we can turn our attention to it if we wish; it remains resolutely present and unified when our active attention is focused elsewhere. The unity of the

> background seems wholly independent of active introspection; it is something active introspection can reveal, when it is appropriately directed, but it is there anyway.[36]

The flaw with this dismissal of the I-thesis (again, as with Tye's argument), is that it does not take into account the difference between *introspectible*, and *introspected*. The I-thesis states that two experiences are phenomenally unified if and only if they are jointly introspectible, without the commitment to the stronger thesis that two experiences are phenomenally unified if and only if they are actually jointly *introspected*. Thus, the I-thesis is not affected by Dainton's above complaint, as the I-thesis does not deny that our background experiences can be phenomenally unified despite not currently being jointly introspected, since all that is required is that they be jointly introspectible. This problem highlights that Dainton does not do enough to keep the notions of introspected and introspectible apart in his discussions of this version of the I-thesis.

Dainton's other arguments against the I-thesis rest for the most part on a notion of the attention-dependence of phenomenal character that he outlines. Recall that any experiences that are phenomenally unified *could be* jointly introspected, according to the I-thesis. This means, Dainton claims, that there will be experiences which we did not in fact introspect, but which we must have been able to introspect, and this can only be the case, i.e. the I-thesis can only be true, according to Dainton, if introspecting or attending to an experience does not alter its phenomenal character, since the identity of an experience depends at least in part on its phenomenal character, at least on one commonplace way of thinking about perceptual experiences.[37] So, if introspection does alter the phenomenal character of an experience, then it is not the case, Dainton claims, that there are experiences which we did not in fact introspect, but could have done. As an example of the kind of attention-dependence that Dainton is talking about, he describes a case in which introspecting a headache can turn it from a low-level background throb into a full-blown horror, and by dint of this, Dainton claims, the phenomenal character is altered.

> Suppose that five minutes ago I was not actively introspecting my slight headache. According to the version of the Weak I-thesis, I could have been attentively aware of this sensation if I had chosen to be. But if I had so chosen my headache would very

[36] Dainton (2006) p. 35
[37] I am aware that this is a potentially controversial statement, but do not have the space to address the issue here.

> probably have intensified, and so the pain I would have been reflectively aware of would not have been numerically the same pain as I actually had.[38]

Pain experiences then, are such that their phenomenal character is attention-dependent, and changes when we attend to the experiences. Other perceptual experiences, Dainton supposes, are similarly dependent: as another example, he consider an auditory experience of a loud babble of voices at a party. By introspecting, separate conversations become recognisable, and the phenomenal character of the experience changes. Since, Dainton claims, it is the case that the phenomenal character of our perceptual experiences does exhibit this attention-dependence, it is not the case that the experiences we have which we did not in fact introspect could be introspected, and thus the I-thesis is false.

There are two responses one could make to Dainton here; the first would be to deny the existence of attention-dependence of the type Dainton highlights, and the second would be to allow that such attention-dependence exists, but deny that this refutes the I-thesis. Dainton claims that introspecting one's headache, to take his example, brings it out of the phenomenal background, and into the foreground, and in doing so, alters the experience's phenomenal character. On the first line of response to Dainton, one could accept that introspecting your experience brings it out of the background, but still deny that in doing this, introspection alters the phenomenal character of the headache experience: we are in fact simply attending more closely to the headache. It seems reasonable to suggest that this is what happens when we introspect our perceptual experiences: the experience is brought to the fore, and the phenomenal character of that experience is more closely attended to, without there being any change in it. It might also be pointed out that on Dainton's line above, there would be no such thing as *introspectible* experiences, as there would be no experiences that we could have introspected but as a matter of fact did not. This seems like a decidedly controversial conclusion to draw, and would require more argument in its favour than Dainton provides in support of the kind of attention-dependence that he posits.

The second line of response to Dainton would be to concede that there is some kind of attention-dependence phenomenon along the lines Dainton describes, and that introspecting the headache experience for example may change the overall phenomenal character of the totality of the subject's overall phenomenal state, but to hold that this

[38] Dainton (2006) p. 38

does not falsify the I-thesis, and hold further, as I will go on to discuss, that this kind of attention-dependence is required for the I-theses to be able to offer any explanation of the difference that phenomenal unity is supposed to make to two unified experiences. What it means to be phenomenally unified is for two states to be jointly introspectible, and if the phenomenal character is altered, then that is just a result of phenomenal unity.

To further illustrate the response that attention-dependence need only alter the subject's overall phenomenal state not the character of individual experiences as they are introspected, consider the following example: there is a headache experience, E_1, which at time t^1 is not currently introspected. At time t^2 the headache is introspected, so that at this time there is E_1, the headache experience, and also I-E_1, which is the introspective state which has E_1 as its object. Here at t^2 there is a change in the totality of the subject's mental states, as both E_1 and I-E_1 are present. There is also a change in the subject's overall phenomenal state, as both E_1 and I-E_1 involve some phenomenal character.[39] We can either say that this is all that happens when we introspect our phenomenal states, or we can push this example further, as follows: at time t^1 again, we have the headache experience that is not being introspected, and at time t^2 there is E_1 and I-E_1. Now, at t^3, after the subject has introspected their headache, Dainton would presumably want to say that there is a different headache experience, E_2, caused by the introspecting of E_1, and thus the headache experience at t^3 is not numerically identical with that at t^1. Alternatively, we might say that at t^3, things are actually as they were at t^1, with the headache experience, E_1, not being introspected. Thus if it is open to us to take this line, then it is open to us to say that although there is some change in phenomenal character caused by introspection, it is only a change in the subject's overall phenomenal state at t^2, and it does not mean that once introspected, the a new headache experience is created which is not identical with E_1.

How then are we to adjudicate here? If Dainton is committed to the idea that introspection alters the phenomenal character of perceptual experiences, and there is reason on the other hand to think that it does not, or that even if it does, that this is not the end of the I-thesis, how are we to settle this? It may seem as though we have something of an impasse here. The situation is complicated further by the fact that though Dainton claims that the truth of the I-thesis cannot be consistent with attention-dependence, in actual fact, if the I-thesis is to serve as a reductive explanation of phenomenal unity, in the way I have discussed, then it may have to rely on some kind of

[39] Though as we have seen from the introduction to this chapter, this might not be the case on all causal accounts of introspection.

attention-dependence to explain the difference that phenomenal unity makes to two unified states. If two states are phenomenally unified, then there is something it is like to have these two states simultaneously, that is different from having them separately. This difference can be explained in one of four ways, as I suggested in the introduction to this thesis: it may be a difference in content, character, both content and character, or some other way that does not involve a change in content or character. A reductive explanation of phenomenal unity in terms of introspective unity will attempt to explain the difference in content or character of the phenomenally unified states entirely in terms of introspective unity, and in order for there to be a change in content or character of the introspectively unified (hence phenomenally unified) states, there needs to be some kind of attention-dependence to explain this difference that phenomenal unity makes. For this reason then, the I-thesis may need some kind of attention-dependence to hold, in order that it might fully explain phenomenal unity. If the I-thesis is to be true, and phenomenal unity just is introspective unity, then introspection will have to alter the introspected experiences, so as to explain the central feature of phenomenal unity. As the potential reductive account of phenomenal unity in terms of spatial unity that was covered in the previous chapter sought to explain phenomenal unity and the difference made by it by positing this difference as that made to two perceptual experiences when they are spatially unified, so this reductive explanation must hold that when two perceptual experiences are introspectively unified, there is difference made to them, and this difference is what phenomenal unity amounts to.

I will return to this issue subsequently, in a later chapter, but first I will examine some further criticisms of Dainton, made by Cody Gilmore.[40] Gilmore offers an appraisal of Dainton's arguments against the I-thesis, and concludes that they fail, before offering an argument of his own to disprove the I-thesis.

Gilmore first considers what he calls Dainton's 'headache argument', which is the argument based on attention-dependence that I discussed above. Gilmore has the following to say about Dainton's position on the I-thesis and attention-dependence (here, I will continue to deal with the one version of the I-thesis, given what I have said above about the difference between the strong and weak versions):

> Consider Dainton's headache experience e and suppose that it is co-conscious with some other experience e*: e could not have been introspected, and a fortiori e and e* could

[40] Gilmore (2003)

> not have been jointly introspected. So e and e* are co-conscious despite failing to be jointly introspectible, hence the I-thesis is false.[41]

Unlike an objection to Dainton which focuses on the existence or not of the kind of attention-dependence which he believes alters the phenomenal character of our perceptual experiences, Gilmore instead focuses on the move Dainton makes from the following claim:

> 1) If it had been the case that I was engaged in the relevant sort of headache-directed introspection time t, then it would not have been the case that the object of this introspection was e (but rather e1, which has been altered by attention, and is thus a numerically distinct experience).

To this claim

> (2) Therefore, it could not have been the case both that I was engaged in the relevant sort of headache-directed introspection at time t, and the object of this introspection was experience e.

Gilmore claims that we can formulate several hypotheses that entail (1), whilst also entailing the negation of (2). The first of theses is

> [H1] the principle that introspection affects the phenomenal character of its objects is a *counterfactual supporting* but *contingent* truth, much like a law of nature. Since this principle is counterfactual supporting, it is true that if Dainton *had* been engaged in headache-directed introspection at t, his headache *would* have had a different phenomenal character. And since the principle is *contingent*, there are possible worlds in which introspection does *not* affect the character of its objects. In some of these worlds, Dainton can and does engage in headache-directed introspection at t without thereby intensifying his headache sensations, hence without preventing his actual headache, e, from occurring. In some such world, e occurs and is introspected.[42]

[41] Gilmore (2003)
[42] Gilmore (2003)

Gilmore claims that since Dainton has given us no reason to believe that introspection *necessarily* affects the phenomenal character of its objects (Gilmore here is rejecting Dainton's attention-dependence), then we have no reason to reject the principle H1 above, and thus no reason to accept premise (2) above, and thus reject the I-thesis. Also, even if Dainton had given us reason enough to accept that introspection necessarily affects the phenomenal character of its objects, Gilmore holds that this would not rule out the following:

> [H2] it is a necessary truth that introspection affects the phenomenal character of its objects. In the actual world, w, Dainton had a highly stressful morning, and his psychological tension built throughout the afternoon. By 6:00 p.m. (time t), Dainton had a headache, e. on a scale of 1 to 10, 10 being the most intense, e was a 'level 3' headache; and, at 6:00 p.m., Dainton was *not* introspectively aware of e. however, in nearby possible world w*, on the very same day, things turned out somewhat differently. Dainton's day was less stressful and his headache sensations were less severe throughout the afternoon. At 5:58 p.m., Dainton had a 'level 2' headache of which he was not introspectively aware. At 5:59 p.m., Dainton happened to turn his introspective attention toward his headache sensations, with the result that they intensified somewhat. At 6:00 p.m., Dainton had a 'level 3' headache, e, of which he *was* introspectively aware.[43]

This, according to Gilmore, also entails claim (1) above, whilst denying (2). It does so because in the actual world, at time t, Dainton was not introspectively aware of his headache, e. This entails (1), and it negates (2) because on the nearby possible world at time t, Dainton was introspectively aware of headache e. So, Gilmore claims, since Dainton has not done enough to show why either H1 or H2 are false, his headache argument fails to undermine the I-thesis. The thrust of these arguments of Gilmore's are that Dainton has not shown us that introspection *necessarily* alters the phenomenal character of our perceptual experiences, and if attention-dependence exists merely contingently, as Gilmore suggests, then this should not be enough for us to reject the I-thesis.

Gilmore's next dissects what he calls Dainton's Shrub Argument, which concerns Dainton's question as to what is responsible for the unity of an introspected experience with a non-introspected experience, with which it is nevertheless phenomenally unified. Dainton believes that it cannot be introspection that is responsible, for the relation of

[43] Gilmore (2003)

unity is not something that is in principle a potential object of introspection. When we introspect part of our current overall experience, and focus on our visual experience of a shrub, this visual experience continues to be phenomenally unified with the rest of your current experiences, and according to Dainton, it cannot be any form of introspection that is responsible for this unity.

> Suppose you do introspect some part of your current experience, actively or passively. This introspected experience remains co-conscious [phenomenally unified] with the remainder of your experience, or at least a large part of it… Given this, what is responsible for the unity of the introspected experience with the non-introspected experiences with which it is co-conscious? One thing seems certain: it cannot be any form of introspection. When you focus your active attention onto the shrub your thoughts, bodily feelings and auditory experiences all remain co-conscious with your visual experience. These experiential relationships cannot be explained in terms of introspectibility, for they are not even potential objects of introspection. If you were to try to actively introspect these relationships you would have to stop introspecting your experience of the shrub.[44]

Gilmore takes pains to turn this into a premise by premise argument, which he stakes out as follows

(1) 'Cav' is an instance of co-consciousness [phenomenal unity] that holds between 'Ea', which is a non-introspected auditory experience, and 'Ev', which is an introspected visual experience of a shrub.

(2) Necessarily: for any experiences e and e* and any instance of co-consciousness c, if c holds between e and e*, then: (i) c is an experience in its own right, and (ii) c is an experience that is co-conscious both with e and with e*.

(3) Therefore, by (1) and (2), 'Cav' is an experience that is co-conscious both with 'Ea' and with 'Ev'.

(4) For any instance c of co-consciousness, if c holds between two experiences only one of which is introspected, then, necessarily, c is non-introspected.[45]

(5) Therefore, by (1) and (4), necessarily, 'Cav' is non-introspected; and a fortiori it is necessary that 'Cav' is not jointly introspected together with 'Ea' (or 'Ev'). In other words, 'Cav' and 'Ea' are not jointly introspectible.

[44] Dainton (2006) p.36-7

[45] Gilmore presumably means to say here 'non-introspectible', rather than 'non-introspected', and I will take him to mean this from here on.

(6) Therefore, (1), (3) and (5), 'Cav' and 'Ea' are co-conscious experiences despite failing to be jointly introspectible; and consequently the I-thesis is false.[46]

Gilmore claims that accepting the truth of (1) and (2), the premise that should be questioned is (4), which states that if c is an instance of phenomenal unity that holds between an introspected experience and a non-introspected experience, then c itself is not introspectible (note: this should be introspectible, rather than introspected as Gilmore has it above).

Dainton's initial defence of this premise is to say that if you were to try and introspect, you would have to curtail your current introspecting of your visual experience. This seems to commit Dainton to the view that it is only possible to introspect one experience at a time. This seems a strange commitment to make without any supporting argument, on the face of it seems obviously false, and Gilmore rejects the idea. However, it is not clear here that Gilmore successfully refutes the I-thesis with this argument. Dainton's argument against the I-thesis here is that introspection cannot be responsible for phenomenal unity, as Dainton claims that the phenomenal unity relation is not something that is even in principle introspectible. Apart from this seeming to presuppose Dainton's own view of phenomenal unity, this does not seem like it hits the target of the I-thesis, as it does not capture how the I-thesis is supposed to work. Rather than introspection somehow revealing the phenomenal unity relation, two perceptual experiences are phenomenally unified *because* they are jointly introspectible, and this joint introspectibility is what gives the two experiences their conjoint phenomenology (which is after all, the feature of phenomenal unity that we are endeavouring to explain).

I will move on now, to look at Gilmore's own argument against the I-thesis, and whether this is something that successfully demonstrates the falsity of the I-thesis. Gilmore bases his argument around the phenomenon of split-brain cases, and holds that his modified example shows that joint introspectibility is insufficient for phenomenal unity. In other words, Gilmore tries to provide an example of two states that are not phenomenally unified, despite being jointly introspectible.

Gilmore's hypothetical twist on the split-brain cases we have looked at involves him having a device which is implanted in his brain, allowing him to voluntarily disconnect the two hemispheres of his brain at will. When the two hemispheres are connected, he supposes, each mental state is phenomenally unified with all other

[46] Gilmore (2003)

experiences in any part of the brain. But when the hemispheres are disconnected, each of them supports its own set of phenomenally unified mental states, so that any state in the right hemisphere is phenomenally unified with all other such states in that hemisphere, and the same goes for the left hemisphere. No two mental states can be phenomenally unified if they are supported by different hemispheres, and the hemispheres are disconnected.[47]

Gilmore then instructs us to imagine a possible world at which the following is true: during t, which is a period of time where Gilmore's hemispheres are disconnected, his right hemisphere supports a visual experience, and his left hemisphere supports an auditory experience. Since the hemispheres are at this point disconnected, the two experiences are not phenomenally unified, nor are they jointly introspectible. However, Gilmore holds that it is plausible to suppose that the two experiences *could have been* jointly introspectible throughout t, since his two hemispheres *could have been* connected throughout t, on some other nearby possible world. He continues, claiming that there is no good reason why it couldn't be the very same experiences that are occurring at this possible world where his hemispheres are connected throughout T, and so, we have a counter-example to the sufficiency clause of the I-thesis.

The problem, however, with Gilmore's argument, is that it seems to make use of a much weakened sense of introspection. In Gilmore's example, the two experiences seem to be introspectible in the same way that any two experiences had by one or even two people are introspectible; two visual experiences, had by two different people, and introspected at the same time, are jointly introspectible, in a very weak sense of introspectible, and it seems that Gilmore is using this sense in his example. It is possible that the two hemispheres of Gilmore's brain *could have been* connected, but it seems equally possible that there is some possible world on which I can introspect my current visual experience, and the visual experience of the person standing next to me, the two experiences being jointly introspectible on this possible world.

As for the above argument of Gilmore's, aside from making use of a very much weakened sense of introspection and introspectibility, there are other concerns with how it is supposed to affect the I-thesis. Gilmore takes his split-brain argument to be a counter to the sufficiency claim of the I-thesis, on which joint-introspectibility is always sufficient for phenomenal unity. His argument refutes this claim, so Gilmore holds, as it shows that we could have an instance where joint introspectibility is not enough for

[47] Gilmore (2003)

phenomenal unity. In order to properly evaluate Gilmore's position, I will take a little time to spell out exactly what Gilmore takes as the I-thesis that he is arguing against.

Gilmore, in elucidating Dainton's position, usefully distinguishes between two possible I-theses that Dainton could be discussing, the *ICC-thesis* and the *IP-thesis*. Both of these theses are attempts to explain phenomenal unity in terms of joint-introspectibility, but the ICC-thesis ties this joint introspectibility to the actual cognitive capacities of the creature in question, rather than to the possibility that the creature in question could possibly have two jointly-introspectible perceptual experiences in some possible world, which is the tenet of the IP-thesis. Gilmore outlines the two theses as follows: firstly, the ICC-thesis

> For any mental states S and S*, any possible world w and any time t in w, S and S* are co-conscious at t in w if and only if: (i) S and S* occur at t in w, and (ii) the owner of S and S* in w has a physical or causal-functional make-up at t in w that gives it the cognitive capacity to jointly introspect S and S* at t in w.[48]

This thesis is spelled out in terms of a subject's actual cognitive capacities, and according to Gilmore, is implausible, as it seems intelligible to think that a raccoon could have phenomenally unified states, without having the capacity to jointly introspect them. Presumably Gilmore would hold the same line with respect to human cases where the subject does not have the requisite cognitive capacities, such as the split-brain cases. I will return to this issue shortly when trying to provide alternative versions of the I-thesis which do not have this problem. Gilmore next defines the IP-thesis as follows:

> For any mental states S and S*, any possible world w, and any time t in w, S and S* are co-conscious at t in w if and only if (i) S and S* occur at t in w, and (ii) S and S* are jointly introspectible; i.e., *there is some possible world in which S and S* are jointly introspected.*[49]

This thesis, as opposed to being tied to a subject, or creature's *actual* cognitive capacities, is explicated in terms of there being some possible world in which a subject can jointly introspect their mental states.

One initial worry that might be generated by these two theses, is that the difference between them is not so obvious, given one way of construing what it is for a subject to

[48] Gilmore (2003)
[49] Ibid.

have a cognitive capacity. One aspect of the supposed distinction seems to be the distinction between a subject 'having a cognitive capacity to jointly introspect S and S* at t in w', which is a condition of the ICC-thesis, and 'S and S* being jointly introspectible; i.e. there is some possible world in which S and S* are jointly introspected' which is part of the IP-thesis. However, the difference between these two conditions might be questioned, as one way to spell out what it is to have a cognitive capacity to φ at world w, is just to φ in some nearby world. If we were to spell things out in this way, it might look as though the two conditions above were equivalent. In response to this worry, we might point to Gilmore's formulation of the ICC-thesis, which would seem to suggest that he does not hold the above conception of what it is to have a capacity: the ICC-thesis states that the creature which is the bearer of the states has to have the requisite physical or causal-functional make-up in this world, not some nearby possible world. Of two states occurring in world w, for these two states to be phenomenally unified, the bearer needs to have the capacity to jointly introspect them in w, and this relies on their having the relevant make-up *in w*, not in some possible world.

Moving on from this worry, with respect to discussion of the I-thesis, Gilmore takes it that Dainton has in mind the IP-thesis throughout his discussion of the I-thesis, but I am not sure that this is in fact the case: Dainton does not discuss the I-thesis in terms of there being some possible world on which the subject's two experiences are jointly introspectible. Besides this, the IP-thesis seems so liberal as to perhaps be trivially true, as we might say that there is no way to rule out there being a possible world on which the subject's states are jointly introspectible, and so the IP-thesis will always be satisfied. This however does not tell us much about what is responsible for phenomenal unity when things are as they are on this world.

To return now to Gilmore's split-brain argument, the target of this argument is the IP-thesis, but as well as attempting to cast doubt on his argument against this thesis, I will also suggest that this version should not be the I-thesis we should be concerned with. Gilmore, like Dainton and Tye, is attacking a version of the I-thesis that should not be put forward in the first place, and so despite showing that these implausible versions fail, they have not shown that there is no plausible version of the I-thesis.

To recap the details of Gilmore's split-brain case: we are supposing that he has a device implanted in his brain enabling him to disconnect his hemispheres at the flick of a switch, and that during t, which is a period of disconnectedness, his right hemisphere supports, inter alia, a visual experience, and his left, an auditory experience. Since these

two experiences, the visual and the auditory, are supported by separate, disconnected hemispheres, they will be both phenomenally disunified, and not jointly introspectible. However, Gilmore argues, there is a sense in which these two experiences *could have been* jointly introspected, since there is a possible world on which the two hemispheres are connected throughout t, and so this scenario would refute the IP-thesis, as it would be a case where two experiences were jointly introspectible, and yet not phenomenally unified.

> [I]t is plausible to suppose that [experiences] eA and eV* could have been jointly introspected, since their respective hemispheres *could have been connected* throughout T, the period when eV and eA occur. That is to say, it seems likely that there is a possible world w* at which my hemispheres are connected throughout T and at which there is a single higher-order state S that introspectively represents both eV and eA. (And if it is not possible that all of this be true of *me,* then surely it is possible that it all be true of *some* sentient being. Let us waive this complication.) There is no good reason to deny that *these very experiences* occur at some possible world w* where my hemispheres are connected throughout T. This is true even if each (token) mental state is so 'modally fragile' that it has its spatiotemporal location, its physical basis, its phenomenal character, and its owner *essentially,* i.e., in every possible world in which it occurs. After all, in light of the set-up of the case, it is clear that all these features can remain constant from possible world w (where my hemispheres are disconnected throughout T), to a possible world w* where my hemispheres are connected throughout T. Nor is there any good reason to deny that at *some* such world where my hemispheres are connected throughout T, the relevant experiences are jointly introspected.[50]

Thus Gilmore takes this to be a counterexample to the claim that joint introspectibility is sufficient for phenomenal unity (a different tack to Dainton, who attacks the necessity claim of the I-thesis).

The problem with this argument of Gilmore's, aside from what has been pointed out above, concerning the very weak sense of 'introspectible' that he is appealing to, is that it is not clear what would prevent the supporter of the I-thesis from simply stating that on this possible world where Gilmore's two hemispheres were connected, and the two experiences were jointly introspectible, that they would also be phenomenally unified. They may simply deny the conceptual distinction between phenomenal unity and joint introspectibility that Gilmore is making. Of course, Gilmore may press the point,

[50] Gilmore (2003)

and argue that there could be *some* possible world on which two experiences could be jointly introspectible, without being phenomenally unified, but it is not clear what then follows from this for the I-thesis. Still, the defender of the I-thesis may respond that though it may be metaphysically (or conceptually) possible for there to be a case of joint introspectibility without phenomenal unity, it is not nomologically possible, and that given that we are interested in what is responsible for phenomenal unity given the laws of nature being as they are, we should not be concerned by Gilmore's example.

Relatedly, we can see why it is not in actual fact the IP-thesis that we should be interested in. The problem is that the IP-thesis is too liberal, and would allow any two states that were had simultaneously by any two subjects to be phenomenally unified, as we could always imagine a possible world in which those two states were had by the same subject.

Thus, an argument against the IP-thesis is not the right approach to take, as this is not the version of the I-thesis that we should be concerned with, being as it is far too liberal and allowing *any* two states to be phenomenally unified. Gilmore is perhaps correct to say that we should not be concerned either with the ICC-thesis, as it may be too restrictive, but there is a potential intermediate position for the I-theorist to take, and I will explicate this thesis when I return to the I-thesis in a subsequent chapter, in the hope of reviving it.

We have seen then, that despite providing more useful criticism of Dainton's treatment of the I-thesis, Gilmore does not himself manage to refute it beyond all doubt. We should think of some version of the I-thesis (though admittedly not the IP-thesis or the ICC-thesis) as still being currently on the table.

6.6 Interim Section

We are now at a point where we can draw some interim conclusions about the material discussed so far in this chapter. Concerning the I-thesis, it seems as though it is still on the table at this point despite its problems, as none of the objections we have seen so far have been overly damaging. I will consider two ways in which we may modify the I-thesis shortly so as to avoid the problems of it being to liberal or too restrictive. Vis a vis Dainton's dismissal of the I-thesis, the arguments that Dainton uses, against the version he outlines, seem unsatisfactory, and there are also ways of explicating an introspective unity thesis such that the issue is not held hostage by the truth or falsity of Dainton's

claims about the attention-dependence of the phenomenal character of our perceptual experiences. Further, Gilmore's objections, particularly those invoking the split-brain cases, which I have looked at above, do not seem to tell conclusively against the possibility of explaining phenomenal unity in terms of joint introspectibility, and again, Gilmore outlines versions of the I-thesis which are not those on which we should concentrate, the ICC-thesis being as he rightly says too restrictive, and the IP-thesis too liberal.

So, we have seen so far in this chapter two versions of the I-thesis that we don't want to accept the IP-thesis and the ICC-thesis, and have also dismissed the criticisms of the general I-thesis from Tye and Dainton. Further, we have seen what the desiderata for an I-thesis would be, and there are two ways in which we might modify the I-thesis, either by adopting a perceptual model of introspection, as outlined at the opening of this chapter, or, by modifying the ICC-thesis. Considering these two options will take up the remainder of this chapter. At the close of this chapter, we will see that despite these two options, the I-thesis will finally have to be abandoned in favour of another explanation of phenomenal unity.

6.7 A Perceptual I-thesis.

As mentioned above, one way to re-draw the I-thesis so as to avoid some of the problems we have seen so far in this chapter would be to make use of the perceptual model of introspection. The motivation behind this move would be that by adopting a perceptual model of introspection, it would be possible to avoid an explanation of phenomenal unity that would seem to rule it out in the case of children and animals, due to the seeming requirement for a sophisticated capacity for introspection. Thus an I-thesis based on a perceptual or inner-sense view of introspection would better meet one of the desiderata for an I-thesis.

What would this new I-thesis look like then? Before we outline the details of this PI-thesis (as it will be known as from here on), it will be worthwhile to re-cap the details of a perceptual or inner-sense theory of introspection.

The central idea behind the perceptual and inner-sense views of introspection is that introspection involves some kind of internal scanning mechanism of our mental states, via some kind of causal process. This can be thought to be analogous to

perception to a greater or lesser degree. So, what commitments need to be made to this kind of view for the PI-thesis to get off and running?

The task is to avoid this PI-thesis falling foul of the same problem that besets the ICC-thesis, namely the commitment to a cognitive capacity that it seems children and animals will not have. The PI-thesis simply needs introspection to involve some capacity for forming a third state which is appropriately related to the two introspected states and which is not so cognitively sophisticated that we would not attribute this capacity to children or animals. This means a commitment to introspection involving some kind of internal scanning mechanism, there being some causal relation between the introspective state and the introspected state, and the introspective state that is formed being perception-like in some important respect.

The PI-thesis needs a causal account of introspection then, where the introspective state has some features in common with a perceptual state. What will these common features be? Firstly, it will need to be the case that the introspective, and target introspected states are related via the appropriate causal relation, analogous to the causal relation at work in perception. Further, it will need to be the case that the introspective state will have as its contents the contents of the target states.

Given these commitments, the PI-thesis can make a case for introspection not involving some sophisticated faculty for higher-order cognitive states, but rather the functioning of some internal scanning mechanism, which is analogous with perception in terms of the kind of causal relation involved and the kind of state formed. This may allow the PI-theorist to avoid the most damaging problem associated with the ICC-thesis.

The PI-thesis then, will run as follows:

> *PI-thesis*: two experiences of a single subject are phenomenally unified if and only if they are jointly introspectible, where joint introspectibility exhausts phenomenal unity, and where introspection involves the operation of a perception-like mechanism.

As with all the other versions of the I-thesis, the PI-thesis works on joint introspectibility rather than needing the two experiences to actually be introspected, and also holds that the two states being jointly introspectible in this way exhausts phenomenal unity, this being necessary in order for the PI-thesis to meet a nothing over and above clause, and

serve as a reductive explanation of phenomenal unity. As before, this issue of reductive and non-reductive explanations will emerge once again shortly.

So, what are the benefits of the PI-thesis? Firstly, does it meet all of the desiderata for an I-thesis which we have been able to outline after consideration of the other versions of the I-thesis considered in this chapter?

After consideration of the IP-thesis outlined by Gilmore, we have seen that it is not desirable for an I-thesis to be too liberal, and allow that any two experiences, not necessarily had by the same subject at the same time, can be phenomenally unified on the basis that there is *some* possible world in which they are jointly introspectible. On the basis of this then, some amendments to the PI-thesis need to be made, so as to avoid this problem. It must be built into the PI-thesis then that the states in question are states of a single subject (as it has been above).

In contrast, after looking at the ICC-thesis, we have seen that we do not want to rule out phenomenal unity in children and some animals due to their not having the capacity to form higher-order states, and thus not having a sophisticated introspective capacity. This is the main motivation for the PI-thesis, and for the adoption of the perceptual model of introspection outlined above.

In this sense, the PI-thesis can actually be seen as a way of modifying the ICC-thesis, by specifying more clearly what the cognitive capacity in question is, and outlining such a capacity that will not rule out phenomenal unity in those subjects or creatures that do not have the capacity to form higher-order states. Once specified, then the potential problem with the ICC-thesis dissolves, as despite this version of the I-thesis relying on the subject in question having some introspective capacity at this world (as opposed to some possible world, as with the IP-thesis), the capacity in question is not one that children and animals would be unlikely to possess, based as it is on some perception-like mechanism which yields a perception-like states as opposed to a higher-order cognitive state. Thus, when the PI-thesis is spelled out, we can see that it charts a course between the problems associated with the two I-theses spelled out by Gilmore.

We can see then, how the PI-thesis might be constructed to avoid the problems raised in connection with the IP- and ICC-theses. How would it fare with respect to some of the other features of reductive explanations of phenomenal unity that we have been considering?

One point at which the PI-thesis would stand to make good sense of a potential feature of phenomenal unity is the 'extra' phenomenology that it may potentially involve.

As with previous discussions of this issue in this thesis, I am not making any claims about whether or not phenomenal unity does actually involve this extra phenomenal character, but am merely examining how the PI-thesis would do in explaining the root of any such extra phenomenology were it posited. The reason the PI-thesis would seem well-placed to explain any extra phenomenology associated with phenomenal unity stems from the perceptual theory of introspection that it involves. If the introspective state which is formed on the basis of the perceptual states is itself a state suitably perception-like, and as such is a phenomenal state, then this state may serve as the bearer of the extra phenomenology that comes with phenomenal unity.

So, a subject has two perceptual experiences which are jointly introspectible, where this involves their having a capacity to form a perception-like introspective state which has as its contents the contents of the two perceptual experiences. If phenomenal unity involves phenomenal character over and above the character of the two perceptual experiences, then it can be attributed to the introspective state. However, as we have seen, a plausible I-thesis will have to be spelled out in terms of introspectibility, rather than introspection, as it is not plausible that all states which are phenomenally unified are also jointly introspected. The PI-thesis is also formulated in this way, and so given this, it seems that it cannot be the introspective state that is responsible for the extra phenomenology in phenomenal unity, as this introspective state may not necessarily be formed, as the subject's perceptual experiences need only to be jointly introspectible, not jointly introspected. Given this, perhaps the PI-thesis is not in much better a position to explain phenomenal unity's extra phenomenology.

An alternative tactic open to the PI-theorist here, if it is thought that any explanation of phenomenal unity must be able to account for the extra phenomenology, may be to suggest that though jointly introspectibility is responsible for phenomenal unity, phenomenal unity does not always involve extra phenomenology, and indeed only does so when the unified perceptual experiences are actually jointly *introspected*. On this line, two perceptual experiences are phenomenally unified when they are jointly introspectible in the right kind of way, i.e. when the subject can form a third, suitably perception-like state on the basis of the perceptual experiences. The two experiences being phenomenally unified does not itself involve any phenomenology over and above the phenomenal character of the two unified experiences, but extra phenomenal character is produced if the two experiences are as a matter of fact jointly introspected,

and thus the bearer of this extra phenomenology is the perception-like introspective state.

This seems the most sensible line to take for the PI-theorist who wishes to endorse the idea that phenomenal unity can involve extra phenomenal character. The difference between the PI-thesis and a view such as Bayne's mereological account, with respect to this question would be that the PI-thesis would not result in the extra phenomenal character being present in all cases of phenomenal unity, whereas Bayne's view, on which two experiences are phenomenally unified when they are actually part of some encompassing state, would mean extra phenomenal character in all cases of phenomenal unity.[51]

As with other occasions where I have brought up this questions, I am not here making any commitments to whether or not phenomenal unity does involve extra phenomenal character, I am merely exploring how the PI-thesis would be placed to explain things if it did.

Moving on to some other points that we need to consider in evaluating the PI-thesis: do the criticisms of I-theses in general made by Tye and Dainton have any purchase against the PI-thesis?

As we have seen above, the main thrust of the criticisms that come from Tye and Dainton are based around phenomenal unity outstripping the experiences that we actually jointly introspect. This is avoided by formulating the various I-theses in terms of introspectibility rather than actually being introspected. As the PI-thesis is also formulated in this way, Tye's and Dainton's criticisms don't affect it.

Another point worth considering is whether or not the PI-thesis would yield a phenomenal unity relation that is transitive, or holds necessarily. In other words, where does the PI-thesis stand in relation to the Unity and Transitivity theses?

As with other versions of the I-thesis, the PI-thesis will give us a phenomenal unity relation that is neither transitive nor holds necessarily. The reason for this, as with other versions of the I-thesis, is that introspectibility does not give us transitivity or the unity thesis. The unity thesis, as we have seen, states that for any two states of a single subject, those two states will necessarily be phenomenally unified. The PI-thesis does not make the unity thesis come out true, as on the PI-thesis, phenomenal unity *is* joint introspectibility (involving a suitably perception-like introspective mechanism) and it is

[51] Of course as I have noted elsewhere in this thesis Bayne does not explicitly commit to phenomenal unity's involving extra phenomenal character. I am merely noting here what his mereological account would involve were such a commitment made.

not the case that for any two states of a single subject those two states will be jointly introspectible. Neither is it the case that the PI-thesis will yield a phenomenal unity relation that is transitive. To see this, consider a subject that has three perceptual experiences at a time, A, B and C. It seems reasonable to suppose that there may be cases where the subject can jointly introspect experiences A and B, and could jointly introspect B and C, but *could not* jointly introspect A and C. Joint introspectibility then, is not a transitive relation, and once again, as phenomenal unity *is* joint introspectibility on the PI-thesis, phenomenal unity will not come out as transitive on the PI-thesis explanation.

As I have previously argued (at greatest length in Chapter Four when discussing the split-brain cases), I do not think that either the unity or the transitivity thesis are theses we should seek to defend, and so I do not see it as a great drawback if on the PI-thesis (and indeed on any variant of the I-thesis) phenomenal unity does not come out as transitive, or as holding necessarily within a single subject.

A more pressing concern for the advocate of the PI-thesis may be whether or not joint introspectibility (of this suitably perception-like kind) is both necessary and sufficient for phenomenal unity, such that coupled with a 'nothing over and above' condition, the PI-thesis can be a genuine reductive explanation of phenomenal unity.[52]

So, is joint introspectibility of the kind outlined in the PI-thesis necessary and sufficient for phenomenal unity? Again, this question has been addressed previously in this chapter with respect to the initial formulation of the I-thesis, and since there seems no good reason to suppose that the answer would change given the PI-thesis advances a particular view of introspection, I will here simply recap what I have said elsewhere in this chapter.

Gilmore takes his modified split-brain case to jeopardise the sufficiency clause of any I-thesis, but as we have seen, not only are there problems with the much weakened use of introspection and introspectibility in Gilmore's argument, he also has as his target here the IP-thesis. As the PI-thesis is actually a modification of the ICC-thesis, the problems associated with the IP-thesis (that it is too liberal as a result of the possible worlds qualification in it) do not apply here. Thus the PI-thesis is in no worse a position here with respect to Gilmore's objection than the ICC-thesis then. Also, there is no reason why the PI-thesis would be any worse off than the original ICC-thesis when responding to putative counterexamples to the necessity claim.

[52] Necessity, sufficiency and the 'nothing over and above' condition are what I am taking to be sufficient for a reductive explanation, as I have explained at the opening of the previous chapter.

One potentially crucial point where the PI-thesis may fail to meet important desiderata however, is in the fact that in advancing the PI-thesis, we must also advance a particular view of introspection. Thus, if the perceptual model of introspection outlined above is implausible, or beset with problems, then the PI-thesis will begin to look less like a plausible explanation of phenomenal unity. Further, not only will the PI-thesis be affected if the perceptual model of introspection is implausible, but it may be though that advancing any I-thesis that necessitates the truth of a particular theory of introspection is inadvisable.

So, just how plausible is the perceptual model of introspection advanced by the PI-thesis? The answer here will depend on just what the commitments of the PI-theorist are: when introducing the PI-thesis, I suggested that the commitments would be that the introspective state, and the relation between this state and the introspected states, be suitably perception-like. With respect to the first of these commitments, I suggested that this would involve the introspective state being a state with content and character, and with respect to the second commitment, that the introspective and introspected states be related by some causal relation. With respect to this second commitment, matters are complicated by disagreement over whether or not perception itself involves some causal relation linking objects in the world to perceptual experiences. This complication can be avoided by the PI-theorist however: they may simply hold that whatever kind of relation links objects and perceptual experiences, an analogous relation links perceptual experiences and introspective states.

Given the motivations behind the PI-thesis, the commitment that needs defence is the first. The PI-thesis is supposed to avoid the issues created by an I-thesis that raises the level of cognitive sophistication required for an introspective capacity so high as to rule out phenomenal unity in some creatures. Given this, the PI-theorist needs to defend the plausibility of the perceptual view of introspection's claims about the introspective state being one which has content and character. The PI-thesis needs this to be true of introspective states in order that the introspective state formed on the basis of two perceptual experiences has the contents of those two experiences as (at least part of) its content, and in order that the introspective state can be one that bears phenomenal character (this being necessary to explain any extra character associated with phenomenal unity). So what is at issue is whether or not introspection could involve the forming of a state with content and character, analogous to a perceptual state.

In the introduction to this chapter I gave a brief outline of the two competing views on introspection, causal views (of which inner-sense views, and the PI-view would be variations) and non-causal views. The objections to the causal and inner-sense views are that a causal relation does not get us what we want from introspection with respect to privileged access, and that the disanalogies between perception and introspection are too great to take seriously any view which makes a close link between the two.

It is with respect to this second complaint that the PI-thesis may fare the worst. Not only does the PI-thesis have to invoke some internal scanning mechanism which links introspected and introspective states, but further, the PI-theorist also has to invoke some introspective mechanism which somehow forms a perception-like state, despite there being no organ of introspection/introspective apparatus. For these reasons, the PI-thesis can be seen as having a serious drawback.

Further, the very fact that the PI-thesis cannot be neutral when it comes to theories of introspection can also be seen as a drawback. Ideally an I-thesis would be neutral as to what particular view of introspection is the correct one.

6.8 A Further Modified I-thesis

Given the shortcomings of the PI-thesis, which centre on the advocating of a controversial view of introspection, is there any alternative way to modify the I-thesis so as to meet the demands we have set?

There is another potential way to modify the ICC-thesis so as to keep the thesis tied to a subject in the actual world (as opposed to the IP-thesis), but not to set the bar for cognitive sophistication too high, nor to make any commitments to a controversial view of introspection (as with the PI-thesis). This approach involves a direct modification of the ICC-thesis set out by Gilmore. As a reminder, the ICC-thesis runs as follows:

> For any mental states S and S*, any possible world w and any time t in w, S and S* are co-conscious at t in w if and only if: (i) S and S* occur at t in w, and (ii) the owner of S and S* in w has a physical or causal-functional make-up at t in w that gives it the cognitive capacity to jointly introspect S and S* at t in w.[53]

[53] Gilmore (2003)

As we have seen, the task here is to provide an explanation of 'cognitive capacity' in this context that is not too restrictive, and as we have seen with the PI-thesis, one way is to build in a perceptual theory of introspection that does not require that the subject in question have the capacity to form higher-order cognitive states. Another approach would be to append the above thesis with a counterfactual claim to the effect that were the subject given the ability to form higher-order mental states, then their physical or causal-functional make-up would be such that they would be able to jointly introspect the relevant states. This modified ICC-thesis would look like this:

> For any mental states S and S*, and possible world w and any time t in w, S and S* are phenomenally unified at t in w if and only if: (i) S and S* occur at t in w, and (ii) the owner of S and S* in w has a physical or causal-functional make-up at t in w that gives it the cognitive capacity to jointly introspect S and S* at t in w, *or*, were the creature to be given the ability to form higher-order mental states, then their physical or causal-functional make-up would allow for a capacity to jointly introspect S and S* at t in w.

Call this the *ICC*-thesis*. If we think that the capacity to jointly introspect two mental states (perceptual experiences or otherwise) requires the ability to form higher-order mental states (and we might think so given our view of introspection), and the ability to form higher-order states is a capacity that children and (some) animals do not have, then we may be able to use the ICC*-thesis as an alternative explanation of phenomenal unity, which still explains it in terms of joint introspectibility, but does not rule out phenomenal unity in children and animals.

Does the ICC*-thesis give the kinds of things desired from an I-thesis? Can it explain any extra character? Based as this version of the I-thesis is on a non-perceptual or non inner-sense view of introspection, it looks as though the ICC*-thesis may not be as well placed to explain this as some other reductive explanations of phenomenal unity. The PI-thesis as we saw above, can potentially explain any extra phenomenology associated with phenomenal unity as being instantiated by the introspective state in instances when it is realised, as this state is a perception-like state with phenomenal character. The S-thesis, which we looked at in the previous chapter, was also well-placed to explain any extra phenomenology, as the extra phenomenology can be explained by the presence of spatial unity, something which carries its own phenomenology.

The ICC*-thesis then, may be at a disadvantage compared to some other reductive explanations of phenomenal unity. Being a higher-order mental state as opposed to a perceptual or perception-like state means the introspective state formed or potentially formed according to the ICC*-thesis is not as obvious a place to locate any extra phenomenology as those given in the examples above.

This of course needn't be disastrous for the ICC*-thesis, as it is not the case that on all views of what phenomenal unity is it must involve extra character, and in most extant descriptions no explicit commitment to extra phenomenology is made. In any case the ICC*-thesis is in no worse a position to explain any extra phenomenal character associated with phenomenal unity than the primitive-relation explanation of Dainton.

As with the other I-theses we have looked at, the ICC*-thesis needs to demonstrate that joint introspectibility is both necessary and sufficient for phenomenal unity, and further that this kind of joint introspectibility exhausts phenomenal unity. I see no reason why the ICC*-thesis will fare differently in this task from the other I-theses, and so will not address this issue again. To the extent that the I-theses in general succeed in showing necessity, sufficiency and exhaustion, the ICC*-thesis will succeed also.

Once again, the criticisms made by Tye and Dainton will not affect the ICC*-thesis, formulated as it is in terms of joint *introspectibility*, and so I will not mention these criticisms further.

The crucial feature of this ICC*-thesis is the counterfactual claim. If the counterfactual added into the thesis makes it implausible, then the ICC*-thesis is in no better a position than the original ICC-thesis, which is too restrictive.

One initial complaint that might be made is that we don't know how to assess the truth of this counterfactual claim, we can't be certain what would happen were the subject in question granted a capacity for higher-order mental states. To diffuse this worry, we might insert a ceteris paribus clause into the ICC*-thesis: the counterfactual would then run '… or were the subject to be given the capacity to form higher-order mental states then, ceteris paribus, their physical or causal functional make-up would allow them to jointly introspect S and S*'.

There are however other objections that might be directed at the ICC*-thesis. The first centres on the potential for counterexamples to the thesis. Certain potential counterexamples will be dealt with by the addition of the ceteris paribus clause, for instance there may be cases where despite the addition of the capacity to form higher-order mental states, the subject still cannot jointly introspect their perceptual experiences.

A further complaint directed at the ICC*-thesis would be that giving a subject who didn't previously have the capacity to form higher-order states that capacity would already be modifying the physical or causal functional make-up of that subject. This complaint would be motivated by scepticism about the possibility of giving a subject a new psychological capacity, without altering the subject's physical or causal-functional make-up.

Does the ICC*-theorist have any response they can make to this complaint? As an initial riposte, they may outline in more detail what would be involved in giving a subject a psychological capacity. To safeguard against the above objection, the ICC*-theorist should hold that what would be taking place would be bringing about the most minimal changes needed to give the subject the capacity for forming higher-order mental states, whilst keeping everything else constant. If this is done, then there are two possible results: the subject can jointly introspect states S and S*, or the subject could not jointly introspect states S and S*. Only the first scenario will result in S and S* being phenomenally unified.

There is an analogy here between this move in the ICC*-thesis, and a view of how we should assess counterfactuals, proposed by David Lewis.[54] On this view, what we are doing in making our assessment of counterfactuals is looking at the closest possible world on which things are that way. Thus what the ICC*-theorist is saying is, go to the closest possible world at which the subject can form higher-order mental states and has no difference in their physical or causal-functional make-up, and if at this world the subject can jointly introspect states S and S*, then these states are phenomenally unified.

Does this successfully defuse the above worry? It may go some way to doing so, but it may steer the ICC*-thesis too close to another problematic issue.

The biggest problem with the ICC*-thesis however, is not the potential counterexamples, or whether or not the idea of giving a subject an additional capacity they previously did not possess, without altering their physical or causal-functional make-up is a plausible one, but the following: the ICC*-thesis attributes phenomenal unity to a subject on the basis of a capacity which they may not have in this world, but at some (albeit close) possible world. This means the ICC*-thesis is still too restrictive, as it still rules out children and animals at *this* world.

The ICC*-thesis then, does not succeed in meeting the desiderata set out for an I-thesis, being as it is too close to the original ICC-thesis in terms of potentially ruling out

[54] Lewis (2001)

phenomenal unity in cases where there may not be any need to. Further, as noted above, it does not seem desirable to have an analysis of phenomenal unity where a subject's phenomenally unified states at this world are to be explained by a capacity that they only possess in some close possible world.

6.9 Conclusion

We are now at a point where we can draw some conclusions about the material discussed in this chapter. Concerning the I-thesis, we have seen several versions. The versions first introduced by Dainton were formulated in terms of two states being introspected rather than introspectible, and as such were vulnerable to the objections raised by Dainton himself and by Tye. Versions of the I-thesis which are formulated in terms of introspectibility however, are immune from these criticisms.

We then saw, via those I-theses introduced by Gilmore that any I-thesis has to navigate between two options, corresponding to different formulations of the conditions for joint introspectibility: being too liberal and allowing any two states to be jointly introspectible, so long as there is *some* possible world on which they are jointly introspectible, or being too restrictive, through demanding that the subject have a cognitive capacity for introspection, where this will rule out phenomenal unity in children and animals. This problem may not even be avoided by adopting a perceptual model of introspection, as with the PI-thesis, as it's not clear if we should attribute even this introspective capacity to children and animals, or indeed to anyone at all.

I then tried to formulate alternative versions of the I-thesis to try and avoid these two extremes. The PI-thesis was an attempt to avoid being too restrictive by adopting a theory of introspection that does not require the capacity for forming higher-order mental states. Despite its advantages, the PI-thesis suffers from two major problems. The first is that it is contingent on the truth of a controversial view of introspection. The second is that even if the perceptual theory of introspection were correct, this might not be enough if it were the case that children and animals still did not possess the capacity for introspection, whether it involves a perception-like state or not. The PI-thesis then, cannot be a successful explanation of phenomenal unity.

Next we saw another attempt to formulate a successful I-thesis, the ICC*-thesis. This was a direct modification of Gilmore's ICC-thesis, where the modification took the form of a counterfactual claim to the effect that were we to give the subject in question the relevant capacity, then their physical or causal-functional make-up would allow them to jointly introspect their perceptual experiences and so enjoy phenomenal unity. The failure of the ICC*-thesis however, lies with this counterfactual claim: in order to make this claim as plausible and to avoid the kind of objections we have seen above, the counterfactual needs to be analysed in terms of the closest possible world at which the subject has the relevant capacity, and everything else is as it is at this world. This however, steers the ICC*-thesis too close to the IP-thesis, and the problem associated with it, i.e. that the ICC*-thesis attributes phenomenal unity to a subject on the basis of a capacity that they don't have in this world, but only in some (albeit close) possible world.

The answer, I will go on to suggest in the next chapter, is to drop the commitment to explaining phenomenal unity using some kind of introspection or joint introspectibility. An explanation of phenomenal unity that does not involve introspection or introspectibility but nevertheless retains the advantages of the views we have seen in this chapter, will be my answer to the question of how to explain phenomenal unity. This will involve dropping the commitment to some kind of I-thesis, but retaining the idea that phenomenal unity is explained by the subject having a capacity to form a third state which has as (at least part of) its contents the contents of the perceptual states. This will in fact steer my view closer to the mereological view of Tim Bayne's that was have seen elsewhere in this thesis, but with some important differences, which I will go into in more detail in the final chapter.

Chapter Seven – A Third-State Explanation of Phenomenal Unity

7.1 Introduction

As we saw in the previous chapter, an attempt to reduce phenomenal unity to introspective unity is rejected by Dainton and Tye, amongst others. Whilst I have conceded above that Dainton and others may be right to reject the I-theses as they formulate them, I also argued that for various reasons, the versions of the I-thesis formulated by Tye, Dainton and Gilmore were not the best possible versions of an I-thesis. Much of the discussion of the I-thesis in the previous chapter is muddied by a conflation, on the part of Tye and Dainton, of introspected and introspectible. Further, Dainton's rejection of the I-thesis is contingent on the existence or otherwise of the kind of attention-dependence that Dainton discusses, and the versions of the I-thesis rejected by Gilmore are either too liberal or too restrictive. Thus, even though Dainton et al take themselves to have refuted the I-thesis, there are alternative versions of the I-thesis which I examined. However, I also showed that even these alternative versions of the I-thesis are beset with problems, and so should not be adopted. At the close of the previous chapter I suggested that it may be commitment to introspection being involved that steers the I-theses into trouble, and there may be alternative explanations of phenomenal unity that can capture the advantages of the I-theses and avoid the drawbacks.

We also saw in the previous chapter what the desiderata for such a modified I-thesis would be: we did not want to exclude children and (some) non-human animals from having phenomenally unified perceptual experiences, nor do we want phenomenal unity to hold between subjects.

These desiderata also hold more generally for any potential explanation of phenomenal unity: it must not be overly restrictive, ruling out phenomenal unity in cases where we would think that it is present. It must also not be overly liberal, positing phenomenal unity in cases where we would think it is absent. These desiderata can both be met by the explanation I will propose in this final chapter.

The form of this concluding chapter will go as follows: in the next section I will present a formulation of a final revised I-thesis. I will suggest that this thesis would work

much more successfully if the commitment to introspection was dropped, and we refrained from invoking a specific mechanism by which the unifying state is formed. This will lead me to suggest that the best explanation of phenomenal unity will be what I will call a 'third-state' explanation of phenomenal unity. The next section of this chapter will be an outline of this third state explanation of phenomenal unity and its relation to other explanations of phenomenal unity. I will then detail how the third-state explanation stands with respect to the desiderata that have been set out, to the transitivity and unity theses, and to other pertinent points that I have looked at in this thesis in relation to other explanations of phenomenal unity. The next section of this chapter will be an attempt to further motivate the adoption of a third-state view, considering also potential lines of objection. I will also here show why this third-state view is a reductive explanation of phenomenal unity. This section will also involve responses to potentially problematic cases for the third-state view, such as split-brain cases and the Sperling experiment. At the close of this chapter, I will claim that a third-state explanation of phenomenal unity is the most successful explanation, and should be adopted.

7.2 One Final I-thesis, The Revised I-thesis

In this section I will introduce one final version of an I-thesis, the intention being to show that this thesis would work much better as an explanation of phenomenal unity if the commitment to introspection of some kind being involved were dropped. The following will serve as a first pass at this revised I-thesis, which will from now on be referred to as the RI-thesis:

> **RI-thesis**: two (or more) perceptual experience are phenomenally unified if and only if they are jointly introspectible in a single act of introspection, where introspection here involves only the ability (on the part of the subject undergoing the experiences) to form a state which draws on the contents of both the target perceptual experiences.

This RI-thesis needs some further unpacking before we can continue however. What needs to be clarified, as mentioned above, is what exactly the notion of introspection being invoked here amounts to and how it relates to the models of introspection outlined at the start of the previous chapter. Further, we would need to set out the following:

what kind of state is formed on the basis of the target experiences, whether or not it is itself a phenomenal state, and whether this introspective capacity invoked in the RI-thesis should be thought of as a capacity of the subject in question, or a property of the token perceptual experiences. In what follows, I will provide some initial defence, or at least details the most sympathetic explanation of this RI-thesis, before pointing out that the thesis is in fact much more plausible without the invocation of a controversial notion of introspection, or introspection at all.

In the RI-thesis stated above, phenomenal unity is explained in terms of the joint introspectibility of the perceptual experiences had by a subject at a time. however, we have seen that it an I-thesis should not involve a sophisticated conception of introspection, as if possession of such an introspective capacity were necessary for phenomenal unity, this may rule out children and animals from having phenomenally unified perceptual experiences, a result which we should be keen to avoid, as I have stated elsewhere.

So given this, what kind of introspection are we talking about when we are discussing the RI-thesis? Fundamentally, this capacity that would be appealed to in the RI-thesis would have to be unlike introspection in many of the ways introspection is standardly conceived. Firstly, the function of this capacity would not primarily be knowledge of our own mental states, as the function of introspection standardly understood is. Relatedly, the state formed when this capacity is exercised is not a higher-order state: it is not a belief or judgement which is 'about' the perceptual state. Relying on a traditional conception of introspection which involves some kind of higher-order state is what steers the I-thesis too close to a thesis which rules out phenomenal unity in animals and children. The state which is formed during an act of introspection as it would be on the RI-thesis is on *the same order* as the perceptual experiences being introspected. It is still a first-order state, which nevertheless draws on the contents of the perceptual experiences had by the subject at that time. In this respect, it would be not unlike the introspective state invoked in the PI-thesis which was outlined in the last chapter. We may even invoke the idea of a maximal perceptual experience here, a notion which has been discussed several times previously in this thesis. There seems no reason to think that a subject's maximal perceptual state at a time is not the kind of state which is being appealed to in the RI-thesis. We might say that a subject enjoys a maximal experience, which has as its contents the contents of the subject's perceptual experiences, just in case they are able to form this state, based on their perceptual experiences, in a

single act. The reason a subject has a maximal experience (which will be phenomenally unified) is that they can jointly introspect all of their perceptual experiences, on the meaning of introspect invoked here.

The notion of introspection invoked in the RI-thesis we can refer to as *non-higher-order introspection*, or *introspection**. Note that the proponent of the RI-thesis does not have to make or endorse any claims about the two kinds of introspection being at odds: the RI-theorist can suppose that an adult human subject will (normally) have the capacity both introspection and introspection*. Again, this is a similar line to one which may be taken by the proponent of the PI-thesis, which also involves a controversial or non-standard notion of introspection. Also similar to the PI-thesis here, is the problem that the RI-thesis may encounter if this introspection* can be shown not to exist.

As a final aside here, one feature that introspection* may have in common with some accounts of introspection, is the kind of incorporation relation that is invoked by Gertler and Chalmers to explain how one state can be (at least) partly constituted by the phenomenal character of its target states. I will not go into this further however, nor will I take any stance here on whether the relation involved in introspection* is a causal or non-causal relation.

Further questions that need addressing whilst spelling out the RI-thesis are 'is the introspective* state a phenomenal state?' and 'is this a capacity of a subject, or is introspectibility* a property of token mental states?'.

Taking the first of these questions, the sensible answer is to say that the introspective* state is a phenomenal state. if we are thinking of he introspective* state as akin to a maximal perceptual experience, formed on the basis of the individual perceptual experiences, one which is (at least) partially constituted by the target experiences. Given that the introspective* state is a phenomenal state, how do we guard against Tye's regress objections against the received view of phenomenal unity?

Recall that Tye questions whether or not the idea of a maximal phenomenal state can be prevented from leading to a regress, since there will need to be something responsible for the phenomenal unity of the individual states with the maximal state, this will need to be a further state, and so on indefinitely. The RI-thesis however, is immune from Tye's regress objection, as it concerns introspectibility. A subject's perceptual experiences at a time are phenomenally unified when and only when they have a capacity to form an introspective* state on the basis of those perceptual state which they have at that time. Phenomenal unity then does not necessitate the maximal state actually being

formed. Even if the state were to be formed, it would be a conjunction of the perceptual experiences, and would not lead to a regress. As Bayne and others say in reply to Tye, there is nothing incoherent about the idea of a maximal perceptual experience.

What about the second question above? Is this introspective* capacity a property of the subject, or is joint introspectibility* a property of the subject's perceptual experiences? The best way to think of introspection* is as a capacity of the subject to introspect* particular token experiences. Though it is a general capacity of the subject, the subject's ability to exercise it may not always be realised, either due to limits on this capacity (such as in the Sperling case) or due to some impairment or problem (such as in split-brain cases).

It should now be clearer what exactly is involved in the RI-thesis, and how it differs from the various versions of the I-thesis discussed in the previous chapter.

Despite attempting to give a sympathetic outline of the RI-thesis, if we consider the following points, we can see that the thesis works more successfully if it does not involve any kind of introspection.

To re-cap, the RI-thesis states that any two states of a single subject are phenomenally unified if and only if they are jointly introspectible in a single act of introspection, where introspection here involves only the ability (on the part of the subject undergoing the experiences) to form a state which draws on the contents of both the target perceptual experiences. As we have also seen above, this is a deflationary notion of introspection, and the RI-theorist (just as with the PI-theorist) may face a struggle in order to motivate acceptance of this kind of introspection, either in place of or alongside more standard notions of introspection.

The motivation however, for the RI-thesis should be clear: the deflationary notion of introspection avoids the RI-thesis being overly restrictive. However, this same desiderata can be met by dropping any mention of introspection from the explanation of phenomenal unity, and instead phenomenal unity is explained by two or more states' being potential parts of a third state, without any commitment to this third state being an introspective or higher-order state.

7.3 A Third-State Explanation of Phenomenal Unity

This 'third-state' explanation of phenomenal unity is an alternative explanation to the various I-theses we have seen, and in this section I will give an initial outline of what such an explanation would look like. As a first description of this third-state thesis, we can say that on this explanation of phenomenal unity, two states of a single subject are phenomenally unified if and only if they are both potential parts of the same third state.

This obviously needs some unpacking. The big question is what kind of state is, and by what kind of mechanism it is formed. Here, the third-state explanation of phenomenal unity can take a cue from the mereological explanation of phenomenal unity put forward by Tim Bayne, which we have looked at previously in this thesis.

As a reminder, Bayne's mereological account holds that two states are phenomenally unified when they are co-subsumed, where this means that they are both parts of a single subsuming state. The subsuming state is taken by Bayne to be a subject's overall phenomenal state at a time, of which the unified experiences are parts.[1] This idea can be borrowed by the third-state thesis, with the modification that the states are unified in virtue of being potential parts of the same overall phenomenal state. The important differences between the third-state and mereological views will be detailed later in this chapter.

As a second formulation of the third-state thesis then:

> **Third-State Thesis**: two (or more) perceptual experiences of a single subject are phenomenally unified if and only if they are both potential parts of the same overall phenomenal state.

Though I will go into more detail later, the most important difference between my third-state view and Bayne's mereological view as he describes it, is that in order to be phenomenally unified on my view, perceptual experiences need only be *potential* parts of the same overall state. Other differences between the two views, notably in relation to the unity and transitivity theses, will be discussed shortly.

Before that, I will here give some consideration of how well the third-state thesis meets the desiderata previously laid out, and how well it may explain certain features of phenomenal unity.

[1] Bayne (2010) Chpt. 2

As I hope to have established, the desiderata that any explanation of phenomenal unity needs to meet are to be neither too restrictive, nor too liberal, and also to successfully account for all the accepted features of phenomenal unity. This last point is crucial for putatively reductive views of phenomenal unity, taking reductive as I am here to be necessity, sufficiency and exhaustion, or a nothing over and above clause.

So, does the third-state thesis successfully avoid ruling out phenomenal unity in cases where we would attribute it, and avoid attributing it in cases where we would rule it out? To begin with, we can consider this issue in relation to the two cases discussed in the previous chapter: phenomenal unity in children and animals, and phenomenal unity across subjects. Given that there is no faculty of introspection at work in the third-state thesis on which phenomenal unity is dependent, then the issue of children and animals is avoided. Likewise, the issue of phenomenal unity across subjects is avoided if we build into the definition of the third-state thesis that it applies only to two or more states of a single subject.

However, these aren't the only cases where phenomenal unity might be controversially attributed or ruled out. As we have seen, split-brain cases are puzzling for studies of phenomenal unity, as it is not entirely clear if we should attribute such unity in these cases. Does the third-state thesis make any attributions here? I will return to this issue in the following sections, after more examination of the features of phenomenal unity that the third-state thesis would have to capture.

Again, I will devote the next section of this chapter to related issues here concerning the transitivity thesis and the unity thesis, so here I will limit myself to other features attributed to phenomenal unity. As I hope to have shown in the first section of this thesis, I think we should consider phenomenal unity to be a relation between token experiences, and the third-state thesis is formulated with this in mind. One potential feature I have looked at with respect to the other explanations of phenomenal unity is the extra phenomenology or phenomenal character it may involve. How would the third-state view explain this? Initially, it may look as though the third-state thesis could simply locate any extra phenomenal character that is to be attributed to the overall phenomenal state. However, since the third-state thesis, unlike the mereological view, is formulated in terms of states being *potential* parts of an overall phenomenal state, this will not work. Instead, the third-state theorist should I think adopt the following view: since it is not part of any received view of phenomenal unity that it must involve extra phenomenology, the third-state theorist should hold (much in the way I suggested when

discussing the PI-thesis) that phenomenal unity does not *necessarily* involve extra phenomenology, but in the instances when it does, i.e. the instances when the unified experiences are actually parts of the overall phenomenal state, the extra phenomenal character is to be located with the overall state. This seems a plausible way for the third-state view to account for extra phenomenal character posited with phenomenal unity, as in most cases it seems safe to assume that the unified experiences will be parts of the overall state, despite this not being the explanation for their being phenomenally unified.

7.4 The Third-State Thesis and the Unity and Transitivity Theses

We have seen in previous chapters of this thesis that on certain views of phenomenal unity, it is taken to be a relation that is transitive and holds within a single subject necessarily. These are both features attributed to phenomenal unity by Bayne, and also by Bayne and Chalmers.[2]

As we have also seen, I do not think that either the unity or the transitivity thesis have been successfully established, and thus we should not attribute transitivity to the phenomenal unity relation, nor should we think that it holds necessarily within a single subject. There are counterexamples to transitivity which I do not think can be fully dealt with, and if we cannot be certain about transitivity, we cannot endorse the unity thesis. Given this, I will not be overly concerned with showing that the third-state thesis sits well with the unity and transitivity theses, as I do not think that either of these theses should be adopted.

Taking this line marks out another difference between the third-state view and Bayne's mereological view. Though there is nothing to say that the mereological view necessarily comes with commitments to the transitivity and unity thesis, Bayne himself endorses both of these, as we have seen. Aside from helping to distinguish my third-state view from Bayne's mereological view, not adopting the unity and transitivity theses also has its advantages when it comes to dealing with problem cases such as the split-brain cases.

In my previous discussions of split-brain cases elsewhere in this thesis, I have concluded that we cannot rule out that there are some split-brain subjects whose consciousness is only partially phenomenally unified. This constitutes a counterexample to the transitivity thesis, and so as we have seen, defenders of the transitivity thesis must

[2] See Chapter Three of this thesis, Bayne (2010) and Bayne & Chalmers (2003).

attempt to show that partial unity is not what is happening in these cases. Further, defenders of the unity thesis must show that in fact there are no breakdowns in phenomenal unity at all in the split-brain cases, and so must provide an alternative explanation that rules out phenomenal unity breakdowns in all split-brain cases.

I on the other hand have suggested that we should not look to impose a blanket view of what is going on in all split-brain cases, and as I am not endorsing either the unity or the transitivity theses, I have no need to do this. Even in advocating the third-state explanation of phenomenal unity, I can make the same claim: since I am not holding the unity thesis, I allow that phenomenal unity can break down within a single subject, and split-brain cases seem good examples of this. If two states are phenomenally unified due to being potential parts of the same overall phenomenal state, then the split-brain cases may seem to be good examples of phenomenal unity breaking down, as the states in question in these cases will precisely not be potential parts of the same overall phenomenal state. Thus split-brain cases are potentially cases of a breakdown in phenomenal unity, but this is not a problem for the third-state view, unattached as it is to the unity and transitivity theses.

7.5 Further Motivation for the Third-State View

We have seen that the third-state view can meet the desiderata for an explanation of phenomenal unity, with respect to not miss-attributing it, and we have also seen how it could potentially explain any extra phenomenology associated with phenomenal unity. What further good reason might we have for adopting the third-state view? Well, we have seen above that when not complicated by commitments to transitivity, the third-state thesis can offer a simple explanation of the split-brain cases. There are other advantages to the third-state thesis that provide motivation for adopting it: as opposed to the primitive relation view of Dainton, the third-state view offers more by way of explanation of phenomenal unity, along the lines of the other explanations we have looked at. Further, and importantly, the third-state thesis seems well poised to explain the phenomenal unity of more than just perceptual experiences. I have throughout this thesis restricted myself to discussion of the phenomenal unity of perceptual experiences, so as to avoid the controversial issue of what other kind of mental states we should also think of as being phenomenally conscious. However, whatever kinds of states turn out to be phenomenally conscious, the phenomenal unity of such states can be explained by the

third-state thesis. As long as they are potential parts of the same overall phenomenal state, then they are phenomenally unified, regardless of whether they are perceptual experiences or beliefs. In this regard, the third-state thesis has an advantage over explanations such as the spatial unity thesis, which as we saw during discussion of it, would struggle to explain the phenomenal unity of states with no spatial content, and as such struggle to explain the phenomenal unity of states other than perceptual experiences.

Now, though explaining the phenomenal unity of states other than perceptual experiences has not been my task in this thesis, it would seem like a general advantage if an explanation of the phenomenal unity of perceptual experiences could be rolled out to explain phenomenal unity more generally also.

This ability to explain phenomenal unity more generally is not unique to the third-state thesis, but does provide some motivation for considering it seriously, especially as other explanations are beset with other problems.

7.6 The Third-State Thesis and Reductive Explanation

One question that emerges with respect to the third-state thesis is in what sense it is a reductive explanation of phenomenal unity.

The question of reductive or non-reductive explanations has arisen at various points in this thesis, and for the sake of clarity, I will briefly go over the ground I have covered again here.

In this thesis I have made the distinction between those explanations which seek to reduce phenomenal unity to some other form of unity, such as spatial or introspective unity, and those that do not. The latter views include those of Dainton, Bayne & Chalmers and Tim Bayne. This distinction has been mostly for dialectical purposes, and I have labelled this distinction a distinction between reductive explanations (the S- and I-theses) and non-reductive explanations (the views of Dainton, Bayne & Chalmers, and Bayne). However, as I have explained, of these latter views, only Dainton's view can be called non-reductive in a stricter sense, seeking as it does to explain phenomenal unity in terms of a primitive relation. All the other explanations we have seen offer some form of reduction, and this third-state thesis is the same, as phenomenal unity is reduced to the states in question being potential parts of an overall state.

As I have also stated previously, when talking about reduction in the strict sense (as opposed to the dialectical distinction above between reductive and non-reductive views), I am taking reduction to involve necessity, sufficiency, and the satisfaction of a nothing over and above clause. Thus for the third-state thesis to be a genuinely reductive explanation of phenomenal unity, being the potential part of the same overall state must be both necessary and sufficient for phenomenal unity, and further it must exhaust phenomenal unity, so that phenomenal unity involves nothing over and above the states in question being potential parts of an overall phenomenal state.

If we take the last of these conditions first, we have already seen that the third-state view can satisfy the nothing over and above condition, there being no features associated with phenomenal unity that the third-state view cannot explain. If phenomenal unity involves extra phenomenal character (and once again, it is never explicitly stated that it does, which in itself takes some pressure off the third-state thesis), then the third-state thesis has a way of explaining this, thus the nothing over and above clause is not violated.

Can the necessity and sufficiency claims be satisfied also? In order to do so, we need to examine the cases which might be seen as counterexamples to the third-state thesis, inasmuch as they may seem to involve phenomenal unity without the unified states being the potential parts of an overall phenomenal state, or states which are potential parts of an overall phenomenal state, but without being phenomenally unified.

In response to other explanations of phenomenal unity, split-brain cases have been touted as counterexamples. Might they be so here also? Split-brain cases might be used to try and jeopardise the third-state view's claims of states being potential parts of the same overall state is necessary for phenomenal unity. They may be used in this way as it would seem that the classic 'key-ring' split-brain experiments reveal states which are not potential parts of the same overall phenomenal state, yet we have also seen claims that these cases are not ones where phenomenal unity is necessarily absent. Bayne & Chalmers and Bayne make the claim that split-brain cases need not involve any breakdown in phenomenal unity, as we have seen.

What should the third-state theorist say in response here? Well they could simply deny that split-brain cases involve any phenomenal unity. There are certainly problems with both Bayne & Chalmers' and Bayne's arguments, as we have seen in an earlier chapter. However, as I have also said earlier, I believe the best treatment of split-brain cases is one in which a blanket interpretation is not imposed on them all, as invariably some cases will not fit well with this interpretation. There will be a large number of split-

brain cases which I would claim involve either full or partial phenomenal disunity, and these pose no threat to the third-state thesis.

The line for the third-state theorist to hold here is that a subject's set of phenomenal states are fully phenomenally unified to the extent that they are all potential parts of the same overall phenomenal state, but being the potential part of an overall phenomenal state remains necessary for any phenomenal unity at all. With respect to split-brain cases, though we should refrain from imposing blanket interpretations on these cases as a whole, it will I believe be possible to show that to the extent that there is any phenomenal unity in a given split-brain case, there will also be potential joint-parthood of an overall phenomenal state of those states which we think are phenomenally unified.

If the third-state explanation can deal with cases such as the split-brain cases, then it will be on the way to securing the claim that potential joint-parthood of an overall phenomenal state is necessary for phenomenal unity. Other cases which would test the necessity claim of the third-state thesis would be dealt with similarly.

One potentially problematic case, which might seem to test the third-state thesis' necessity claim involves the Sperling experiment, which was discussed in Chapter Four. To recap, in the Sperling experiment subjects were presented with a matrix of letters, which was flashed on a screen for 250ms, then a tone sounded indicating to the subject whether they were to report the top, middle, or bottom row of the matrix, or the whole matrix. The subjects reported seeing the whole matrix when it was flashed in front of them, but could only report the contents of one row. As we saw, this leads Bayne & Chalmers (2003) to claim that the Sperling experiment is an instance of phenomenal unity, but access disunity, as they also classify split-brain cases. The Sperling case may present the following problem for the third-state thesis: given that the subjects can only accurately report the contents of one row of the matrix, but also report seeing the whole matrix, this might be a case of a phenomenally unified perceptual experience, in the absence of joint potential parthood of an overall phenomenal state. If this were the case, it would cast doubt on the necessity clause of the third-state thesis.

We have seen why this kind of case causes problems for the various I-theses, as it would seem to show that phenomenal unity outstrips joint introspectibility of any kind. We have also seen that the responses available to the I-theorist are not all that plausible: An extreme response may be to discount the subject's response that they had a visual experience of the whole matrix, and hold that they do not have a phenomenally unified

experience of the whole matrix. This response would mean that the Sperling case is no longer a threat to the various I-theses, as there would not be phenomenal unity without joint introspectibility. A more concessive response may be to say that in the case of the Sperling experiment, the subject has a single visual experience, all of the contents of which are not introspectible, though some are. They can form an introspective state which has some of the contents of the visual experience, and would thus have a degree of phenomenal unity. This second response would have the result that perceptual experience can outstrip phenomenal unity, and we can have a perceptual experience which is only partially phenomenally unified. Though this might sound implausible, if we allow, as I am, that cases of partial phenomenal unity are possible, then if they can happen between perceptual experiences, then why not within experiences also?

Another tack when responding to the Sperling experiment case is to challenge the current orthodox position which is as follows: cases like the Sperling experiment show that phenomenology outstrips access. This orthodoxy is the foundation for Bayne & Chalmers' claims about the Sperling case.

This orthodox account of the Sperling experiment has been challenged however. In recent work, Ian Phillips has argued that there are better alternative interpretations of the Sperling experiment which do not support any claims about phenomenology outstripping access.[3]

Though this last approach may be a fruitful one for the defender of the I-thesis, much has still to be done to overturn the orthodox interpretation of the Sperling cases, and if this is not done, then the other responses open to the advocate of the I-thesis do not look appealing.

Unlike the various I-theses however, the third-state thesis and its necessity clause is not jeopardised by the Sperling case. The reason for this is that the third-state thesis does not involve joint introspection or introspectibility, but rather potential joint-parthood of an overall phenomenal state. Thus even though more than four letters of the matrix are not jointly introspectible, they can still be potential parts of the same overall phenomenal state. Indeed, as the subjects report seeing the whole matrix, this may well be the case.

The Sperling experiment shows up the limits of joint attention or joint introspection, but this does not imply anything about being potential parts of an overall phenomenal state. Thus the Sperling experiment does not serve as a counterexample to

[3] Phillips (forthcoming) and Phillips (forthcoming) (b)

the third-state thesis, as it is not a case of phenomenal unity without potential joint-parthood of an overall phenomenal state.

This highlighting of the Sperling experiment and the third-state thesis' explanation of it not only reinforces the third-state explanations claims to being a genuinely reductive explanation of phenomenal unity, it also shows another advantage this explanation of phenomenal unity has over the various I-theses, which struggle to give a plausible response to the Sperling case.

Are there any cases which challenge the sufficiency claim of the third-state thesis? In order to be a successfully reductive explanation of phenomenal unity, as well as satisfying the nothing over and above clause, and being necessary for phenomenal unity, it must also be the case that there are no instances of potential parthood of an overall phenomenal state without phenomenal unity also.

The most immediately obvious examples that might be touted as such cases are those which would seem to involve more than one subject. We can imagine, it might be claimed, cases where two states of two different subjects are nevertheless both potential parts of the same overall phenomenal state. Nevertheless, we should not attribute phenomenal unity in this case. There are several reasons why we might think that states across subjects may still be potential parts of the same overall phenomenal state: we might hold a mereological view of overall phenomenal states, or we may read 'potential' in such a way as to mean that two states are potential parts of the same overall phenomenal state if there is some (not necessarily nomologically) possible world on which these two states are part of the same overall phenomenal state. Here then, potential joint-parthood would not be sufficient for phenomenal unity.

How should the third-state theorist respond to these kinds of cases? The response should I think to deny that the third-state view defines 'potential' in the way it is above, and instead hold that potential is limited to the actual world, so that two states being potential parts of the same overall phenomenal state is true if and only if they are potential parts of the same overall state at this world, and that overall phenomenal states at this world do not bridge across separate subjects. This way we will not get cross-subject cases of joint potential parthood without phenomenal unity. This way, the third-state theorist can defend the sufficiency claim needed to make the third-state explanation a reductive explanation of phenomenal unity.

7.7 Further Thoughts, Objections and Replies

There are several other points worth considering in relation to the third-state thesis: how it stands with respect to the issue of the individuation of experiences, being as it is closely related to Bayne's mereological account. Also, it is prudent to consider some of Tye's regress objections to received views of phenomenal unity, as if the third-state explanation deals with these successfully, this is yet more motivation to adopt it.

As I outlined at the start of this chapter, the third-state explanation of phenomenal unity is close Bayne's mereological view, the crucial difference being that on the third-state view, states need only be *potential* parts of an overall phenomenal state in order to be phenomenally unified, whereas Bayne holds that two states are phenomenally unified when and only when they are co-subsumed.[4]

As Bayne points out, his mereological conception of phenomenal unity does not sit well with certain views of the individuation of experiences. in particular, it is at odds with Tye's one experience view.[5] Without going over this ground in great detail again, it is worth pointing out where the third-state explanation stands here. The third-state view will also be at odds with Tye's one experience view, as it takes phenomenal unity to be a relation that holds between token states, and further, allows that these states can be proper parts of encompassing phenomenal states. I will address Tye's objections again shortly, but will here make a final brief point on the individuation of experiences.

Bayne holds that his mereological view sits well with what he calls the tripartite conception of experiences, outlined elsewhere in this thesis. He also holds, as I do, that there is no uniquely privileged method of individuating experiences, and that the tripartite method is simply his preferred method in the context of issues of phenomenal unity.[6] So, in advocating the third-state explanation of phenomenal unity, am I committing myself to a particular view of the carving up of experiences? I do not think so. The third-state view fits with all views on which phenomenal unity is a relation between token states, and so I need not commit myself any more specifically than that. I will have to reject Tye's one experience view, but I have provided reason elsewhere in this thesis to reject, or at least be suspicious of Tye's view.

[4] Bayne (2010) p. 21
[5] Ibid.
[6] Ibid. p. 24

On the subject of Tye once more, we have seen his objections to what he calls received views of phenomenal unity. Does the third-state explanation fall foul of these objections?

Recall, Tye's objections to received views take the form of two perceived regresses: one a regress of states, one of content. I have outlined the methods used to deal with these regress objections elsewhere in this thesis, so will offer the briefest of recaps here, before suggesting that the third-state explanation of phenomenal unity can avoid these complaints also. The pertinent objection here is Tye's second regress objection, which states that if phenomenal unity relies on a maximal unifying state, then this state must need a further state to unify it with those states which unifies, and so on ad infinitum. The response given by Bayne and others is to say that the unified states are parts of the maximal state, and that parthood is a reflexive relation, therefore the maximal state is part of itself and can be self-unifying in this sense.[7] Further, this applies to the third-state explanation only in the instances where the unified experiences are actually parts of the overall phenomenal state (though this will be the majority of the time). In instances where the unified states remain as potential parts, Tye's regress objection will not get started. The third-state thesis is therefore safe from the objections Tye levels at received views of phenomenal unity.

7.8 Conclusion

We have seen in this chapter then, that there is a viable reductive explanation of phenomenal unity in terms of the perceptual experiences of the subject being potential parts of an overall phenomenal state, and that this potential joint-parthood is necessary and sufficient for phenomenal unity and further phenomenal unity is nothing over and above this potential joint-parthood. This gives us a successful reductive explanation of phenomenal unity, albeit not one that is reductive in the sense of reducing phenomenal unity to some other kind of unity exhibited by consciousness. The third-state explanation performs better than the other explanations considered in this thesis, and also has the advantage of being able to explain the phenomenal unity of all phenomenal states, beyond simply explaining the phenomenal unity of perceptual experiences.

[7] Bayne (2010) p. 29. See also Chapter Two of this thesis for more discussion of the responses available to the received view theorists.

The third-state explanation has the advantage over the views of Bayne and Bayne and Chalmers of not being wedded to the unity and transitivity theses, and so being forced to impose blanket interpretations on empirical cases such as split-brains, which I have argued are best dealt with on a case-by-case basis. More generally, we have seen that despite Tye's objections, we should think of phenomenal unity as a relation that holds between token states, and also between token state parts, and that despite being couched in alternative terminology, Tye's own theory is no alternative view.

The third-state view then offers a plausible reductive explanation of phenomenal unity: potential joint-parthood of an overall phenomenal state is necessary and sufficient for phenomenal unity, and further, there is nothing to phenomenal unity over and above this.

Bibliography

Bayne, T. (2005) 'Divided Brains and Unified Phenomenology: A Review Essay on Michael Tye's Consciousness and Persons' *Philosophical Psychology* 18 (4) 495-512

- (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: Kluwer pp. 201-10.

- (2010) *The Unity of Consciousness* Oxford: Oxford

- (forthcoming) 'Perception and the Reach of Phenomenal Content' in F. Macpherson and K. Hawley (eds.) *The Admissible Contents of Experience* Wiley-Blackwell

Bayne, T. & Chalmers, D. (2003) 'What is the Unity of Consciousness?' in A. Cleeremans (ed.) *The Unity of Consciousness: Binding, Integration and Dissociation* Oxford: Oxford University Press pp. 23-58

Bertelson, P., Francesco, P., Ladavas, E., Vroomen, J. and de Gelder, B. 'Ventriloquism in Patients with Visual Neglect' *Neuropsychologia* 38 1643-42

Block, N. (1995) 'On a Confusion About a Function of Consciousness.' *Behavioural and Brain Sciences* 18:227-47

Bogen, J. E. & Vogel, P.J. (1962) 'Cerebral Commissurotomy in Man: Preliminary Case Report' *Bulletin of the L.A. Neurological Society* 27 169-72

Brewer, B. (1995) 'Bodily Awareness and the Self' In Jose Luis Bermudez, Anthony J. Marcel & Naomi M. Eilan (eds.), *The Body and The Self*. (Cambridge, Mass: MIT Press.)

Burge, T. (1988) 'Individualism and Self-Knowledge' *Journal of Philosophy* 85

Byrne, A (2005) 'Introspection' *Philosophical Topics* 33 (1) 79-104

Campbell, J. (1994) *Past, Space and Self* (Cambridge: MIT Press)

Carruthers, P. (2010) 'Introspection: Divided and Partly Eliminated' *Philosophy & Phenomenological Research* 80 (1) 76-111

Chalmers, D. (2002) 'The Content and Epistemology of Phenomenal Belief' in Q. Smith and A. Jokic (eds.) *Consciousness: New Philosophical Essays* Oxford: Oxford University Press

- (unpublished manuscript) 'Verbal Disputes and Philosophical Progress'

Dainton, B. (2006) *Stream of Consciousness* 2nd. Edition London: Routledge

- (2004) 'Higher-Order Consciousness and Phenomenal Space: reply to Meehan' *PSYCHE* 10 (1)

- (2007) 'Coming Together: the Unity of Conscious Experience' in M Velmans and S Schneider (eds.) *The Blackwell Companion to Consciousness* Oxford: Blackwell pp. 209-23

Dennett, D. (1979) 'Where am I?' in D. Dennett (ed.) *Brainstorms: Philosophical Essays on Mind and Psychology* Montgomery: Bradford Books

Evans, G. (1982) *On the Varieties of Reference* edited by John McDowell Oxford: OUP Press

Gazzaniga, M. S., Bogen, J. E. & Sperry, R. W. (1963) 'Laterality Effects in Somesthesis Following Cerebral Commissurotomy in Man' *Neuropsychologia* 1 209-15

Gertler, B. (2009) 'Introspection' in P. Wilken, T. Bayne & A. Cleeremans (eds.) *The Oxford Companion to Consciousness* Oxford: Oxford University Press pp. 386- 90

Gilmore, C. (2003) 'The Introspectibility Thesis' *PSYCHE* 9 (5) http://www.theassc.org/vol_9_2003

Grice, H. P. (1962) 'Some Remarks About the Senses' in R. J. Butler (ed.) *Analytical Philosophy, First Series* Oxford: Basil Blackwell pp. 133-53

Hurley, S. (1998) *Consciousness in Action* Cambridge: Harvard University Press

Lewis, D. (2001) *Counterfactuals* 2nd. Edition Oxford: Blackwell

Levy, J. (1990) 'Regulation and Generation of Perception in the Asymmetric Brain' in C. Trevarthen (ed.) *Brain Circuits and Functions of the Mind: Essays in Honour of Roger W. Sperry* Cambridge: Cambridge University Press pp. 231-48

Lockwood, M. (1989) *Mind, Brain and the Quantum* Oxford: Basil Blackwell

- (1994) 'Issues of Unity and Objectivity' in C. Peacocke (ed.) *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind* Oxford: Oxford University Press pp. 89-95

Macpherson, F. (2006) 'Ambiguous Figures and the Content of Experience' *Nous* 40 (1) 82-117

- (forthcoming) (ed.) *The Senses: Classical and Contemporary Readings* Oxford: Oxford

Meehan, D. (2003) 'Phenomenal Space and the Unity of Conscious Experience' *PSYCHE* 9 (12) http://www.theassc.org/vol_9_2003

Milner, B., Taylor, L. & Jones-Gotman, M. (1990) 'Lessons from Cerebral Commissurotomy: Auditory Attention, Haptic Memory, and Visual Images in Verbal-Associative-Learning' in C. Trevarthen (ed.) *Brain Circuits and Functions of the Mind: Essays in Honour of Roger W. Sperry* Cambridge: Cambridge University Press pp. 293-303

Milner, B., Taylor, L. & Sperry, R. W. (1968) 'Lateralised Suppression of Dichotically Presented Digits after Commissural Section of Man' *Science* 161 184-6

McGurk, H. &MacDonald, J. (1976) 'Hearing Lips and Seeing Voices' *Nature* vol. 264 pp. 746-748

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

Nudds, M. (2001) 'Experiencing the Production of Sounds' *European Journal of Philosophy* 9 (2) 210-29

- (2010) 'Sounds and Space' in M. Nudds and C. O'Callaghan (eds.) *Sounds and Perception: New Philosophical Essays* Oxford: Oxford pp. 69-97

O'Dea, J. (2008) 'Transparency and the Unity of Experience' in E. Wright (ed.) *The Case for Qualia* Cambridge: MIT Press

Phillips, I. (forthcoming) 'Attention and Iconic Memory' in C. Mole, D. Smithies and W. Wu (eds.) *Attention: Philosophical and Psychological Essays* Oxford: Oxford University Press

- (forthcoming) (b) 'Perception and Iconic Memory' *Mind and Language*

Raymont, P. and Brook, A. (2009) 'Unity of Consciousness' in B. McLaughlin, A. Beckermann & S. Walter (eds.) *The Oxford Handbook of Philosophy of Mind* Oxford: Clarendon Press pp. 565-77

Revonsuo, A. (2003) 'The Contents of Phenomenal Consciousness' *PSYCHE* 9 (12) http://www.theassc.org/vol_9_2003

Richardson, L. (2010) 'Seeing Empty Space' *European Journal of Philosophy* 18 (2) 227-43

Schechter, E. (2010) 'Individuating Mental Tokens: the Split-Brain Case' *Philosophia* 38 195-216

Searle, J. (2000) 'Consciousness' *Annual Review of Neuroscience* 23 557-78

Seigel, S. (2010) 'The Contents of Perception' *Stanford Encyclopaedia of Philosophy* www.plato.stanford.edu

Sorensen, R. (2008) Seeing *Dark Things: The Philosophy of Shadows* Oxford: Oxford University Press

Sperling, G (1960) 'The Information Available in Brief Visual Presentations' *Psychological Monographs* 74/11

Sperry, R. W. (1974) 'Lateral Specialisation in the Surgically Separated Hemispheres' in F. O. Schmitt & F. G. Worden (eds.) *Neuroscience, 3rd Study Programme* Cambridge: MIT Press pp. 5-19

Strawson, P. F. (1959) *Individuals* London: Methuen

Teng, E. L. & Sperry, R. W. (1973) 'Interhemispheric Interaction During Simultaneous Bilateral Presentation of Letters or Digits in Commissurotomised Patients' *Neuropsychologia* 11 131-40

Tye, M. (2002) 'Representationalism and the Transparency of Experience' *Nous* 36 (1) 137-51

- (2003) *Consciousness and Persons: Unity and Identity* Cambridge: MIT Press

Wilson, D. H., Reeves, A. G., Gazzaniga, M. S. & Culver, C. (1977) 'Cerebral Commissurotomy for control of Intractable Seizures' *Neurology* 27 708-15

- (1982) 'Central Commissurotomy for Intractable Generalised Epilepsy, Series Two' *Neurology* 32 687-97