

2016-03-12

The Spatial Structure of Unified Consciousness

Bartłomiej Chomanski
University of Miami, b.chomanski@gmail.com

Follow this and additional works at: https://scholarlyrepository.miami.edu/oa_dissertations

Recommended Citation

Chomanski, Bartłomiej, "The Spatial Structure of Unified Consciousness" (2016). *Open Access Dissertations*. 1580.
https://scholarlyrepository.miami.edu/oa_dissertations/1580

This Embargoed is brought to you for free and open access by the Electronic Theses and Dissertations at Scholarly Repository. It has been accepted for inclusion in Open Access Dissertations by an authorized administrator of Scholarly Repository. For more information, please contact repository.library@miami.edu.

UNIVERSITY OF MIAMI

THE SPATIAL STRUCTURE OF UNIFIED CONSCIOUSNESS

By

Bartłomiej Chomanski

A DISSERTATION

Submitted to the Faculty
of the University of Miami
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Coral Gables, Florida

May 2016

©2016
Bartłomiej Chomanski
All Rights Reserved

UNIVERSITY OF MIAMI

A dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

THE SPATIAL STRUCTURE OF UNIFIED CONSCIOUSNESS

Bartłomiej Chomanski

Approved:

Elijah Chudnoff, Ph.D.
Professor of Philosophy

Berit Brogaard, Ph.D.
Professor of Philosophy

Mark Rowlands, Ph.D.
Professor of Philosophy

Guillermo Prado, Ph.D.
Dean of the Graduate School

Tim Bayne, Ph.D.
Professor of Philosophy
University of Manchester and Western University

CHOMANSKI, BARTLOMIEJ
The Spatial Structure of Unified Consciousness

(Ph.D., Philosophy)
(May 2016)

Abstract of a dissertation at the University of Miami.

Dissertation supervised by Professor Elijah Chudnoff
No. of pages in text. (138)

In this dissertation, I defend two theses: first, that experiences with spatial phenomenology represent space as single and unitary, and, second, that every experience has a spatial phenomenology. The two theses entail the conclusion that the unity of conscious spatial representation is necessary for the unity of consciousness. This means that unified consciousness has a (partially) spatial structure. In what follows, I adopt the mereological conception of unified consciousness according to which experiences are unified just in case they are parts of a subject's total phenomenal state. I also assume that the unity relation is a relation among experiences.

Most philosophers working on the unity of consciousness reject the conclusion that unity of consciousness requires the unity of spatial representation. But in my dissertation I show that there are compelling reasons to accept the two theses that jointly entail it. First, I argue that all spatial perceptual representations represent space as single and unitary. Then, I argue that non-perceptual spatial experiences: imagination experiences, recollective experiences, as well as experiences of afterimages and phosphenes have the kind of spatial content that relates the locations represented in these experiences to the locations represented in perception. In short, the locations represented in perceptual and non-perceptual experiences are represented as belonging to the same space.

Secondly, I argue that experiences standardly taken not to possess spatial phenomenology, such as conscious thoughts and moods, either do in fact possess such phenomenology (thoughts), or wholly depend for their own phenomenology on spatially unified experiences (moods).

When it comes to conscious thoughts, I argue that the view that they exhibit spatial phenomenology is dialectically privileged over its denial. This means that granting spatial phenomenology to conscious thoughts is a more reasonable default starting position. Considerations marshalled from the philosophical accounts concerning what it's like to have conscious thoughts, as well as those marshalled from the literature on the pathologies of cognitive experience (especially schizophrenia) strongly suggest that what it's like to have thoughts does involve awareness of the thoughts' locations (typically in one's head).

When it comes to mood experiences, I argue that all it takes to account for the phenomenal character of such experiences are modifications to other, non-mood experiences. Such modifications are not, however, free-standing experiences for which the question of unity arises. This is because the unity relation is a relation between experiences, not between non-detachable aspects of a single, free-standing experience. Hence, the view that for the mood experience to obtain, other experiences must simply be modified in certain ways is compatible with the view that spatial unity is necessary for phenomenal unity. This is because the relation between an experience and its modification is not the phenomenal unity relation. The way mood experiences are unified with other experiences is wholly captured by the way in which a modification of an experience is unified with that experience. And this is not a phenomenal unity relation. On the other hand, the experiences

that collectively give rise to the mood experience are spatially and phenomenally unified.

Hence, mood experiences do not pose a special problem for my account.

I conclude that unified consciousness requires unified conscious spatial representation. Consciousness has a spatial structure.

Moim Rodzicom, Lidii i Janowi.

And to Anna.

Acknowledgments

I was fortunate to work on this dissertation with some of the most amazing and supportive philosophers I've ever met. I am extremely grateful for this opportunity to the entire philosophy department at UM (with special thanks to Professors Amie Thomasson and Brad Cokelet).

In particular, I'd like to thank my committee: Professors Elijah Chudnoff, Berit Brogaard, Mark Rowlands, and Tim Bayne. Thanks to you all for guiding me through the various stages of writing and thinking about these things.

Thank you, Tim, for your incisive and careful comments on the previous versions of this work. Your work on consciousness has been inspirational for my thinking about it.

Thank you, Mark, for your sharp, challenging comments always delivered in the kindest possible way. They definitely helped me improve this work in many ways.

Thank you, Brit, for your helpful and encouraging comments and suggestions, and for bringing to bear your encyclopedic knowledge of the literature on this work.

And, finally, thank you, Eli, for being such an amazing committee chair. Your encouragement, your keen eye for problems with my arguments and opportunities to develop them, and your willingness to guide and support me throughout working on this project have been invaluable. I cannot emphasize enough how much your mentoring has helped me at every stage of writing and thinking about philosophical problems. Thank you.

I also wish to thank other graduate students in the department with whom I discussed philosophy (and everything else) and from whom I learned a great deal. Adam, Ben, Ben, Casey, Dan, David, Haley, Fredrik, Jordan, Lance, Mark, Nihel, Noel, Nurbay, Philipp, Rina, Robin, Ryan, Sarah, Ted, Zach; a huge thanks to you guys!

Table of Contents

List of Figures	viii
List of Tables	ix
Chapter 1: Introduction	1
Chapter 2: Unity of Consciousness and the Experience of Space.....	9
2.1 Spatial Unity and Perceptual Experience.....	9
2.2 The Spatiality Thesis and Bayne’s Objection.....	11
2.2.1 “Borgy”	12
2.3 Motivations for the Disunity Judgment	14
2.3.1 Frames of Reference	15
2.3.2 Failure to Meet the Embodiment Constraint	18
2.3.3 No Single Frame	21
2.3.4 No Spatial Relations	26
2.4 Spatial Disunity and Phenomenal Unity are Incompatible.....	27
2.5 Summary.....	36
Chapter 3: Imagination and Memory Experiences are Spatially Unified with Perceptual and Somatic Experiences	38
3.1 Non-Perceptual Spatial Experiences.....	38
3.2 Imagining, Remembering, and Perceiving Space.....	38
3.2.1 Superimposing Images	40

3.2.2	Localizing Imagery.....	43
3.2.3	Perspectivalty	45
3.2.4	Memory Experiences.....	50
3.2.5	Unity without Distance?	51
3.3	Non-spatial Seeing.....	52
3.3.1	Does RM Fail to Perceive Space?	52
3.3.2	Growing Determinacy of Spatial Vision	56
3.4	Seeing without Connecting.....	60
3.5	Summary.....	62
Chapter 4:	A Case for Spatial Cognitive Phenomenology	64
4.1	Do You Experience Your Thoughts as Occurring in Your Head?	64
4.2	Spatial Phenomenology in Conscious Thought	70
4.2.1	Topological Properties.....	71
4.3	Introspective Evidence of Spatial Cognitive Phenomenology.....	74
4.4	Cognitive Phenomenology and Spatial Cognitive Phenomenology	77
4.5	Thought Broadcast and Spatial Cognitive Phenomenology	82
4.5.1	Thought Broadcast.....	82
4.5.2	Spatial Phenomenology in Thought Broadcast	83
4.5.3	A Possible Objection	87
4.6	Summary.....	90

Chapter 5: What makes up a Mood Experience?	91
5.1 Moods and Their Characteristics	91
5.1.1 Dimensions of Mood Theories	92
5.1.2 Classifying Mood Theories	93
5.2 Modification.....	97
5.2.1 Examples of Modification in General	97
5.2.2 Modifications and Moods.....	100
5.2.3 Beyond Modifications	102
5.3 Patterns of Modification	103
5.4 Mood Experiences Depend on Patterned Modifications	105
5.5 Explaining Mood Phenomenology	108
5.5.1 Characteristic Features of Moods.....	108
5.5.2 PMV Explains Everything about Mood Phenomenology	111
5.5.3 Mood Phenomenology and the Competitors to PMVs.....	116
5.5.4 Amy Kind’s Challenge	120
5.6 Mood Experiences and Representationalism.....	122
5.7 Mood Experiences and Spatial Unity of Consciousness.....	122
5.8 Summary.....	125
Chapter 6: Conclusion.....	126
Bibliography	131

List of Figures

- Figure 2.1 This array could be seen as a single scene, with the cloud-shape in the leftmost position, the sun-shape in the rightmost position, and the remaining shapes: a smiley face, a vertical line, a moon, a rectangle, another vertical line and a heart shape in between. But Borgy could see each pair of images separated by lines by means of a different body, without relating them all to a single frame of reference. 21
- Figure 2.2 The images from left to right: a moon-shape, a rectangle, a sun-shape and a cloud-shape 22
- Figure 3.1 The arrows indicate the directional axes, the grey circles indicate objects of imagination, and black squares indicate objects seen. 47

List of Tables

Table 5.1 A summary of how various views account for central aspects of mood phenomenology.....	110
---	-----

Chapter 1: Introduction

As you're giving a lecture on the mind-body problem, you see row after row of captivated listeners, you hear and understand your own voice, you're wondering to what extent your exposition is clear, you feel the pressure the soles of your feet exert on the ground as you realize that the temperature in the classroom could be raised a degree or two. You're a bit overexcited, as you've just had your third coffee this morning, and it has left a slight aftertaste in your mouth. You sense a headache coming, and you're a little hungry.

There's not only something it is like to have each of these experiences individually. Your experience of the hunger and your experience of your own voice are connected in ways that *your* experience of the hunger and your *listeners'* experience of your voice are not. There is nothing it is like for your listeners' experience and your own experience to be had together. There is no "bigger" experience these individual experiences are parts of. But there *is* something *it is like for you to have your experiences together. There's something it is like to be you at that moment.* All the conscious states had together by you are parts of your overall stream of consciousness. They are unified. That unity (i.e. the *phenomenal unity of consciousness*) is the topic of this dissertation.

Let us say that unified consciousness consists of a number of phenomenally conscious states, connected by some relation in virtue of which they have conjoint phenomenology. Let us call this relation the phenomenal unity relation. The phenomenal unity relation accounts for the phenomenal difference between there being a special phenomenal character to my hearing x and seeing y simultaneously, and there being no such phenomenal character to my hearing x and *your* seeing y simultaneously. This is because

my auditory experience and my visual experience stand in the phenomenal unity relation, while my auditory experience and your visual experience don't. Simply put, for my conscious states to be unified is for there to be something it is like to experience them together. But could more be said about phenomenal unity? Could an *account* of phenomenal unity be given?

Let's say that unified consciousness is *structured* just in case an account of the phenomenal unity relation in more basic terms *can* be given. More precisely, phenomenal consciousness is structured just in case necessary and sufficient conditions for conscious experiences to be phenomenally unified can be found.

What are the more basic terms in which the phenomenal unity relation can be accounted for? Here's my proposal: in addition to being phenomenally unified, a subject's conscious states might stand in a number of "coherence relations." Coherence here is understood very broadly: it refers to some kind of integration among either conscious experiences themselves, or the contents of these experiences, along a parameter or a range of parameters. Coherence relations are the more basic terms in which to account for phenomenal unity.

Now, one answer to the question about the structure of consciousness is, obviously yes. Unified consciousness is very often structured in various ways. For instance, all your conscious states right now might happen to have mutually consistent contents. Or, they may all happen to be realized in some single part of the brain. Yet, saying this doesn't seem to give us any insight into the essence of conscious unity. A philosophically more interesting question in the vicinity seems to be whether any kind of structure is necessary and/or sufficient for phenomenal unity.

One answer to this question is that unified consciousness does not have any necessary structure, because phenomenal unity is a primitive relation. This seems to me to be Barry Dainton's (2000) view.

Those who do not take Dainton's route can explain the unity relation in two different ways: either by appealing to the relations between the *contents* of the unified experiences, or by appealing to relations between the experiences themselves.

Tim Bayne (2010) offers a prominent account of phenomenal unity of the latter sort. According to Bayne, consciousness has a mereological structure: two experiences are phenomenally unified just in case they are parts of the subject's total phenomenal experience. Bayne explains this view by, first, endorsing what he calls "the 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." Second, he accounts for the relation of *being subsumed by* in terms of the relation of *being a part of*. This account grants unified consciousness more structure than Dainton's—it endows consciousness with a part-whole structure—but Bayne doesn't seem to think that unified consciousness necessarily displays any more structure than this. In particular, it is not necessary that there be any ordering among the subsumed experiences. We might call the kind of structure granted to consciousness by Bayne's theory, a *thin* structure.

The views proposed by Sebastian Watzl (2014), Luke Roelofs (2014), and Elijah Chudnoff (2013) endow unified consciousness with a "thicker" structure than Bayne's view. On these views, unified consciousness is structured into a center and a periphery (i.e. into some sort of a phenomenal field). According to these theorists, unified experiences, in

addition being parts of the total phenomenal state, also exhibit some extra ordering, in terms of which some conscious states end up in the center, and others end up in the periphery of unified consciousness. Like Bayne's view, however, these views also do not recruit the contents of the partial states to do any essential work in explaining the structure of unified consciousness.

Michael Tye's (2003) is a content-based mirror image of Bayne's view, as far as the structure of consciousness is concerned. For Tye, a subject has a unified consciousness just in case the subject has a single, total phenomenal experience with multiple, multimodal contents. Again, though, no further ordering amid the contents of experience seems necessary on Tye's account. Susan Hurley (1998) offers a similar content-based account of unified consciousness.

A content-based account that grants consciousness more structure than Hurley's and Tye's views is Sydney Shoemaker's (1996) view. According to Shoemaker, "unity of consciousness is in part a matter of one's various beliefs forming, collectively, a unified conception of the world." This account of unity demands more from the contents of unified consciousness than merely being conscious contents entertained simultaneously. Some amount of representational coherence (though what exactly *this* is is hard to explicate) is also required of the individual contents. In particular, some contents might be more central to the picture of the world than others. (In fact, Hurley's view could be placed between Tye's and Shoemaker's, because Hurley also requires that the contents of unified states be subject to some objective norms of coherence; Tye doesn't appear to require that).

Yet another content-based view is proposed by Farid Masrour (2014). On Masrour's view, phenomenal unity is not explicable in terms of any single relation among the contents

of phenomenal states, but rather in terms of many relations that we may be aware of in our experience: causal relations, spatial relations, rational relations, and so on. The structure of consciousness on Masrour's view is heterogeneous.

In developing their own accounts, many theorists of phenomenal unity (Dainton (2000), Bayne & Chalmers (2003), Bayne (2004, 2010), Tye (2003), Watzl (2014), and Roelofs (2014)) consider yet another way in which the phenomenal unity relation can be explained by a content-based view. They all reflect on, and reject, the view according to which *spatial* content might hold the key to a structural account of phenomenal unity. The view in question is that phenomenal unity might be explained in terms of spatial unity of experience.

Bayne and Chalmers offer a helpful (though, as will emerge in Chapter 2, too narrow) characterization of spatial phenomenal unity: "a set of experiences are spatially unified if (i) each has spatial representational content, and (ii) the representational content of each is *comparable*, in the sense that the objects represented are represented as being in spatial relations to each other [*italics in original*]." Accounting for phenomenal unity in terms of spatial unity has several virtues: (1) it allows us to make sense of the typically perspectival nature of unified consciousness (because spatial representation is perspectival); (2) it allows us to capture an aspect of unified consciousness as a picture of a unified world that Shoemaker speaks about, without overintellectualizing unified consciousness (spatial representation is plausibly available for lower animals and infants; building up a *conception* of the world, not so much); (3) it allows us to understand phenomenal unity in terms of representations whose neuro-functional profiles are relatively well-understood; (4) it captures an intuitively appealing naïve introspective judgment; after all, it just seems

that my experiences represent to me a spatially unified world. If explaining phenomenal unity in terms of spatial unity is so great, whence the resistance?

There are two main objections to the idea that phenomenal unity can be explained in terms of spatial unity, denying either condition (i) or (ii) in Bayne and Chalmers's definition. On the one hand, some conscious states seem to lack spatial representational content (e.g. conscious thoughts), or even lack representational content altogether (e.g. moods). On the other hand, actual subjects of experience may enjoy phenomenally unified spatial experiences that are not comparable (e.g. imagining space and perceiving space at the same time); plus, there are possible subjects of experience who could simultaneously perceive different locations that are not comparable, in virtue of having their sensory organs receiving input from spatially distant places at the same time.

I propose that all these objections can be resisted. Hence, in this dissertation, I argue for the view *that spatial unity of consciousness is necessary for the phenomenal unity of consciousness*. I deal with these objections in the reverse order to the one they were just presented. I start with possible and actual putative spatial disunity, and then talk about putative lack of spatiality.

I base my thesis on two premises:

- (1) Every experience is a spatial experience (or wholly depends on spatial experiences).
- (2) All of the subject's phenomenally unified spatial experiences are experiences as of the same, single, unitary space.
- (3) So, all of the subject's phenomenally unified experiences are experiences as of a single unitary space.

I support premise (1) mainly by considerations contained in Chapters 4 and 5, where I talk about experiences that may initially seem to fail to exhibit any spatiality: conscious thoughts and mood experiences, respectively. Some reasons to think that sensory experiences (including olfactory experiences and auditory experiences) and bodily experiences have spatial phenomenology are scattered throughout Chapters 2 and 4. A more sustained discussion of visual, and yet arguably non-spatial, experiences undergone by a patient suffering from Balint's syndrome is contained in Chapter 3. In general, though, I focus on conscious thoughts and moods. This is because it is not phenomenologically apparent that these sorts of experiences exhibit any spatiality at all. On the other hand, auditory, olfactory, and gustatory experiences seem to the introspective eye to be at least *prima facie* spatial.

Here is a slightly more rigorous presentation of considerations supporting premise (1):

- (4) Cognitive phenomenology has a spatial aspect (Chapter 3).
- (5) Sensory phenomenology and bodily phenomenology have a spatial aspect (Chapters 2-4).
- (6) Mood phenomenology is wholly dependent on affective modifications to sensory and cognitive experiences, hence on spatial experiences (Chapter 5).
- (7) So, every experience is a spatial experience, or wholly depends on spatial experiences (=Premise 1).

I support premise (2) by considerations contained primarily in Chapter 2, where I motivate the principle that all of the subject's perceptual experiences represent a single, unitary space. I then devote a section of Chapters 4 and 5 to show how non-perceptual experiences that are the focus of these chapters are spatially unified with perceptual

experiences. Chapter 3, in turn, argues that quasi-perceptual experiences and some atypical perceptual experiences are also spatially unified with perceptual experiences.

Here is a slightly more rigorous presentation of considerations supporting premise (2):

- (8) All phenomenally unified perceptual experiences of a subject are experiences as of the same, single, unitary space (Chapter 2).
- (9) Imagination experiences, memory experiences and phosphene/afterimage experiences of a subject are spatially unified with the subject's perceptual experiences (Chapter 3).
- (10) A subject's experiences of the character that patients with Balint's syndrome have are spatially unified with the subject's other experiences (Chapter 3).
- (11) A subject's cognitive experiences are spatially unified with the subject's perceptual experiences (Chapter 4).
- (12) A subject's mood experiences are wholly dependent on the subject's spatially unified experiences (Chapter 5).
- (13) So, all of a subject's spatial experiences are experiences as of the same, single, unitary space (=Premise 2).

It follows that there cannot be phenomenally unified experiences that aren't also spatially unified. Spatial unity is necessary for phenomenal unity.

I also take it that spatial unity is obviously sufficient for phenomenal unity: to experience any two objects or locations as belonging to the same space, I have to experience them together. This explains why the objections raised to explaining phenomenal unity in terms of spatial unity focus on the necessity component of the account.

Chapter 2: Unity of Consciousness and the Experience of Space

2.1 Spatial Unity and Perceptual Experience

In this chapter I argue that there is a deep connection between the unity of perceptual-somatic consciousness and the experience of space. In particular, I argue that for one's perceptual-somatic consciousness to be unified, one's conscious states have to represent a spatially unified environment. I consider a popular objection to this line of thought, according to which it's possible to have a unified consciousness while simultaneously experiencing spatially disconnected environments. In responding to this argument, I develop an account of spatial experience, and the experience of spatial connectedness that entails the necessity of the latter for phenomenal unity.

There are several questions that can be raised with regards to phenomenal unity of consciousness: phenomenological questions (what is unified consciousness like from the inside?), structural questions (what are the necessary and sufficient conditions for consciousness to be unified?), and explanatory questions (in virtue of what is consciousness unified?), among others (the distinctions and connections between these questions are explored in more detail in (Masrour, 2014), and the seminal paper by Tim Bayne and David Chalmers (2003). While these questions are obviously distinct, one might find it tempting to treat them as interdependent, so that the answer to one constrains or informs answers to the others.

One way to connect the issues raised by some of the above questions would be to say this: having a unified consciousness feels a lot like experiencing a spatially connected world from a point of view. This (answering the phenomenological question), in turn,

suggests a way to pursue the remaining questions; that is, it suggests exploring the relation between spatial experience and phenomenal unity. For starters, take Michael Tye's description of what it's like to have a unified conscious experience:

At the current time, sitting on a patio by a swimming pool with my computer while sipping some tea, I have a single perceptual experience with a rich and multimodal phenomenology. My experience represents to me visual, auditory, olfactory, gustatory, and tactual features in relation to myself. I experience the red flowers to the left of here, the diamond shape of the tiles on the patio relative to me and my viewing position, the sounds of a bird straight ahead, the smell of the chlorine from the pool off to my right, the smoothness of the computer keys against my fingers, the taste of the tea as it comes into contact with my tongue. The world, as I perceptually experience it, thus, is a world in relationship to here (Tye, 2003)

Tye's description captures a pervasive feature of our perceptual engagement with the world. Our experience is suffused with (egocentric) spatial content, i.e. with how things seem to us *from here*. Moreover, all those "features" mentioned by Tye seem to occupy the same, single space. It is thus tempting to wonder if having experiences that represent worldly happenings as occurring in a single space (i.e. having your experience exhibit *spatial* unity) might be a good candidate for being the necessary and/or sufficient condition for *phenomenal* unity. It seems clear from introspective evidence that they go together in many ordinary situations. Do they have to?

While this line of thought is initially tempting (at least as far as the necessity claim is concerned), most philosophers who have considered it ultimately wish to reject it (Bayne, 2010; Bayne & Chalmers, 2003; Dainton, 2000; Roelofs, 2014; Watzl, 2014). One reason to do so presents itself already when addressing the phenomenological question: some conscious states (e.g., conscious thoughts) seem to lack any spatial phenomenology. There are reasons to be skeptical of such an outright rejection and I explore them in Chapters 4 and 5. In this chapter I address a different kind of argument to the effect that spatial unity

is not necessary for phenomenal unity. That is, the context here is the answer to the structural question, and not the phenomenological question. Even if someone is persuaded, however, that spatial phenomenology is absent from some kinds of conscious states, one could take what follows as disputing the more restricted claim that spatial unity is not necessary for the unity of *perceptual* consciousness.

In the next section I discuss a thought experiment that is supposed to undermine the necessary connection between spatial unity and phenomenal unity. I then examine a number of ways to motivate the judgment that spatial unity breaks down in the situation described by the thought experiment. I then criticize each motivation as insufficient to justify the disunity judgment.

2.2 The Spatiality Thesis and Bayne's Objection

In this and the following sections, I examine and criticize an objection to the following way of answering the structural question: consciousness is phenomenally unified just in case it is spatially unified. The objection I will mostly focus on is Bayne's (2010) "Borgy" thought experiment. The thought experiment may be understood as a challenge to the claim that spatial unity is necessary for phenomenal unity (I will call this claim the "spatiality thesis"). This way of reading the thought experiment has been noted in recent writings on the unity of consciousness by Watzl (2014) and Roelofs (2014). Interestingly, Bayne doesn't take the spatiality thesis to be his main target. Watzl and Roelofs, on the other hand, take Bayne's example to have shown the falsity of the spatiality thesis. But it seems to me that Bayne's thought experiment does not go as far as Watzl and Roelofs need it to go in order to cast serious doubt on the spatiality thesis. As a result, the main purpose of this paper is to block a popular line of thought in theorizing about unity of consciousness,

according to which Bayne's thought experiment is sufficient to establish the falsity of the spatiality thesis. While doing so, I sketch a (hopefully) plausible account of spatial unity and spatial experience that entails the spatiality thesis.

As a side note, it's worth mentioning that Bayne's thought experiment is a variation on a theme pursued by Tye and Dainton (cf. his (2000)). The latter two philosophers have constructed cases where a single subject's sense organs become spatially separated (vision in one place, bodily sensations and proprioception in another), so that the subject consciously represents different spaces at the same time. I focus on Bayne's variant as it sidesteps (as Bayne himself notes) issues concerning the potentially necessary connection between somatic experience and sensory experience that Dainton's and Tye's thought experiments might run afoul of. I think, however, that the points I raise about Bayne's take on the thought experiment apply equally to Tye's and Dainton's versions.

2.2.1 "Borgy"

Here's Bayne's thought experiment:

Consider an organism who I shall call 'Borgy'. Borgy is a 'scattered creature', for he has three bodies that are biologically independent of each other. (We might think of each of these bodies on the model of the human body.) However, the brains that are located in these three bodies can communicate with each other via miniature radio transmitters. In fact, Borgy's brains are functionally integrated with each other to roughly the same degree that our two neural hemispheres are. Although Borgy has three bodies, he is a single subject of thought, action, and experience. His perceptual states feed into a unitary cognitive system, and he has direct control over each of his three bodies (this, in part, is what makes the three bodies *his* bodies) ... The fact that he has three bodies allows Borgy to experience the world from the perspective of three non-contiguous locations at once. At one and the same time he might feel the sun beat down in Sydney, hear a plane flying over the Sonoran desert, and smell freshly baked croissants on the steps of Sacré Coeur. Although Borgy's bodily and perceptual experiences will be structured around distinct bodies, I suggest that they might well be phenomenally unified. Borgy's three brains are distributed between his three bodies, but I see no reason to deny that they might jointly implement the kind of neurophysiological singularity that underlies the unity

of consciousness. This ‘sub-personal’ integration might also have a ‘personal’ echo. Because Borgy's perceptual experiences feed into a single cognitive system, he will be able to experience certain relations between events that occur within his individual perceptual fields. *Even though he doesn't experience the spatial relations between his three bodies*, his phenomenal perspective could include a representation of (say) the relationship between the temperature in each of the locations in which he has a body ... To take another example, suppose that there is a duck in front of each of Borgy's three bodies. *Borgy will not experience the three ducks as spatially related to each other*, but he may be able to experience one of the ducks as bigger ... than the others. The upshot of these considerations, I suggest, is that Borgy constitutes a counterexample to the embodiment constraint, for it is plausible to suppose that he could retain a phenomenally unified consciousness despite the fact that neither his perceptual experiences nor his bodily sensations will be structured around a single physical object. [my emphases] (2010)

Bayne's aim, as can be seen from the above quotation, is to cast doubt on what he calls “the embodiment constraint.” The embodiment constraint says that for one to have a phenomenally unified conscious experience, one's perceptions and bodily sensations need to be structured around a single body. Borgy has a phenomenally unified experience without it being structured around a single body (or so it seems); so, the embodiment constraint is false. I take no issue with this line of thought in what follows. But Watzl and Roelofs take the Borgy thought experiment to prove more than the inadequacy of the embodiment constraint. They take it to show that the spatiality thesis is false. Here's one way to see how the above thought experiment serves to motivate the denial of the spatiality thesis:

- (1) Borgy's experience is phenomenally unified.
- (2) Borgy's experience is spatially disunified.
- (3) If Borgy's experience is phenomenally unified and spatially disunified, then spatial unity is not necessary for phenomenal unity.
- (4) So, spatial unity is not necessary for phenomenal unity.

Premise 3 above is incontestable. Bayne replies convincingly to worries concerning premise 1. This leaves the proponents of the spatiality thesis with premise 2 to object to. In the remainder of this paper, I try to show that premise 2 is mistaken.

As it stands, premise 2 is supported by brute, untutored intuition. At first sight, it seems obvious that Borgy's experience represents disconnected regions of space. It's not clear to me, however, given what's at stake (i.e. showing that an intuitively attractive starting point for a theory of unified consciousness is wrong), that we should be satisfied with mere intuition.

2.3 Motivations for the Disunity Judgment

In this section, I will examine a number of more theoretically robust motivations for premise 2. I will ask: "Under what conditions are we entitled to say that spatial unity breaks down?" Answering this question will give us a way of getting clearer about the kinds of considerations relevant to assessing the level of threat Borgy poses for the spatiality thesis than what reliance on mere intuition would get us.

Obviously, one wouldn't get far in this enterprise, were one to leave the notion of spatial unity without explication. A general idea would be this: any two of the subject's conscious perceptual states S1 and S2 are spatially unified just in case (a) S1 and S2 both have some spatial representational content, and (b) S1's and S2's spatial representational content is such that the regions represented by S1 and S2 are experienced as parts of/belonging to the same space.

To have spatial representational content is to represent an object, event, or a place as being/occurring somewhere. I leave the notion of experiencing two regions or objects as belonging to the same space at an intuitive level. One clear example of such an experience,

however, would be an experience of two objects (or regions) as standing in some spatial relations, such as left-right or front-back, to one another.

Now that our basic terms are defined, I will present three ways of motivating premise 2. I dub them as follows:

- (a) Failure to Meet the Embodiment Constraint: one route to motivate premise 2 would be to say that Borgy *doesn't experience his environment as centered around one spatially continuous object* and, as a result, his experience is spatially disunified.
- (b) No Single Egocentric Frame: another way to motivate the spatial disunity claim is to say that *the spatial content of Borgy's experiences cannot be structured by a single egocentric frame of reference*. That is, there is no one frame of reference that may be used to specify all of the spatial content of Borgy's experiences.
- (c) No Spatial Relations: yet another is to say that Borgy *represents no spatial relations between objects represented in different streams*.

I think that none of the motivations presented above suffices for the acceptance of premise 2. In the following sections, I outline my reasons for rejecting each.

2.3.1 Frames of Reference

The discussion of some of the abovementioned ways in which one can motivate the thought that Borgy's experience is spatially disunified will rely on the idea of frames of reference *structuring* experience in various ways. We need to spend some time on unpacking the notion of the frame of reference, and what it is for it to structure experience.

There are at least two ways of thinking about frames of reference, depending on how robustly one understands that notion. On the one hand, we might think that frames of reference essentially consist of the origin of the frame and the spatial relations objects and

places in a scene are represented as standing to the frame's origin and to one another. In this vein, Joel Smith writes that "a frame of reference is an origin and set of axes that define a direction system." (Smith, 2014). The axes usually mentioned in this context are front/back, left/right, and above/below. Similarly, Irvin Rock describes a frame of reference as "a unit or organization of units that collectively serve to identify a coordinate system with respect to which certain properties of objects, including the phenomenal self, are gauged" (quoted in (Levinson, 1996). Along the same lines, Bill Brewer and Stephen Pears (1993) talk about frames of reference as devices for individuating places in terms of their spatial relations to some origin.

In contrast, Roberta Klatzky's explanation is of a deflationary kind: "a reference frame is a means of representing the locations of entities in space" (1998). However, if we think about frames of reference along the lines of Klatzky's definition we will not be able to avail ourselves of any means of individuating frames of reference. For this reason, I shall interpret "frame of reference" robustly, that is, as involving an origin and a set of axes in terms of which spatial relations between objects can be identified.

What is it for a frame of reference to *structure experience*? We can understand that phrase to mean, roughly, this: an experience E is structured by a frame of reference R just in case what it's like to undergo E (E's phenomenal character) is in part determined by R, which means that it is partly in terms of R that places (locations of objects) and relations between them are consciously represented by E. For example, suppose you see an apple on a table. The apple appears to stand in spatial relations to other items in the scene, e.g. it appears to be in front of and slightly to the left of your torso, above the floor, on top of the table etc. The frame of reference that structures your experience, then, is one that has you

(or some part of your torso) as an origin, and the usual left/right, front/back, above/below axes defining the direction system. Now, imagine you extend your hand in front of you so that your palm occupies the leftmost portion of your visual field. You can still see the apple as to the left, your palm as more to the left, and both of these in front of your torso. But another way of seeing such a scene is to see the apple to the right of your palm, and your torso as behind it. If this was the way the scene would appear to you, we could say that the frame of reference that has your palm as its origin (and the usual left/right etc. axes) structures your experience.

Both the change in the origin and the axes of a structuring frame of reference can make a phenomenal difference to your experience. That difference may be manifested in e.g. experienced sameness of location. Building on the example above, suppose you move your hand so that your palm is now on the opposite side of the apple. If it seems to you that the apple is no longer on the right, then your experience is structured by the frame of reference centered on your palm. If it seems to you that it's still on the left, then the structuring frame is centered on your torso. There is a phenomenal difference between seeing the apple as remaining on the left and seeing it as now on the left and previously on the right. Hence, the difference in the origin of a frame of reference that structures your experience makes a phenomenal difference.

The difference in the type of directions defined by the frame of reference may also affect phenomenology. In the example above, the directions are *relative* to the origin of the reference frame (left/right, front/back). But they needn't be. Your experience could be structured by an absolute frame of reference. E.g. it's possible to feel the wind as blowing from the East, even as you turn your body around, so that you feel the wind first on your

face, then on the left side of your body, then on your back. It seems different than experiencing the wind as blowing first from the front, then from the left, then from the back, etc. Hence, the difference in the kind of axes that define a system of directions in a frame of reference that structures experience can likewise make a phenomenal difference.

Conscious spatial representation is typically a representation of how the world appears in relation to the viewer. That is, it is typically structured by an egocentric frame of reference. An egocentric frame of reference is one whose origin is the subject's point of view.¹ While typically an egocentric frame of reference will have as its origin a region on the subject's body, it's not easy to say what is necessary and sufficient for a frame of reference to count as egocentric. Proposals include the immediate connection to the possibilities for the subject's spatial action (e.g. Campbell, 1993; G. Evans, 1982), the representation of the subject's spatial properties in an indexical way (e.g. Smith, 2014), as well as the identity of the origin of the frame with the subject's body (e.g. Pick Jr & Lockman, 1981).

Without taking a stand on this issue, and borrowing a term from Rock, I want to commit myself merely to the claim that for a frame of reference to count as egocentric, it is necessary that it have the "phenomenal self" as its origin.

2.3.2 Failure to Meet the Embodiment Constraint

One tactic to motivate premise 2 would be to argue that Borgy's experience is spatially disunified because the egocentric frames that structure it are centered on distinct objects

¹ Perhaps this definition is so non-committal as to be completely uninformative. But any further assumption about what makes a given frame of reference egocentric is likely to be controversial. Additionally, "view" need not be taken in its literal sense. When I hear a sound as coming from the left, my experience is still structured by an egocentric frame of reference.

that aren't (non-scattered) parts of any larger object. This criterion allows for ordinary perceptual experience to count as spatially unified. But it fails to rule in an imaginary experience that, it seems to me, would be a good example of a spatially unified experience.

This is why: there does not appear to be any conceptual difficulty in imagining one's bodily sensations (say, of touch) to occur in an object one experiences as spatially disconnected from the rest of one's body. Suppose I'm undergoing a rubber-hand illusion, and the rubber hand is gradually moved further and further away from my body. While I don't know what the actual experience in such an experiment would be like,² it is at least conceivable that I *could* continue to feel sensations in the hand as it moves away from the rest of my body. Now, if I could do that, I could also undergo perceptions with spatial egocentric content specified by the frame of reference centered on the rubber fingers or the rubber palm. And, if I could do that, I could also, it seems, integrate the content of those perceptions with the egocentric spatial content given by the frame of reference centered on the rest of my body. But since I'd be able to do that, I don't think this would entail an experience of two disconnected spaces; that is to say, nothing in my experience would suggest that I'm experiencing two disconnected spaces.

More abstractly, it appears possible to seem to feel bodily sensations in objects that are not physically connected to my own body, and are not experienced as spatially contiguous with my body. And, furthermore, there is nothing incoherent in supposing that we could also undergo conscious perceptions whose spatial content is specified by a frame of

² Actual experiments (Armel & Ramachandran, 2003; Lloyd, 2007) suggest, though, that the illusion weakens as the distance increases; on the other hand, it is still felt in elongated hands, hands at an anatomically impossible distance from the subject, and ordinary objects like tables under the right circumstances.

reference centered on those detached body parts (e.g., I can touch the far side of a long table and experience it as under my fingertips, and at such and such a distance from my torso, despite the rest of my body being at the near side of the table). Now, it seems clear to me, however, that at no point in the rubber hand illusion did I begin to experience two disconnected spaces. And if that's right, then the mere spatial discontinuity between the origins of the egocentric frames of reference does not suffice for spatial disunity.

If the rubber hand illusion seems far-fetched, imagine a different scenario: imagine losing all proprioceptive sensation in your torso. It seems possible that you'd still integrate proprioceptive inputs from your feet, your hands and your head (your experience would be different than normal, but it would not be, I submit, an experience of disconnected spaces) into a conscious representation of a single space. It would still seem to you that you occupy a single space, even though now some parts of your body are experienced as not connected to other parts.

As a result, it seems possible to (a) have an egocentric spatial experience whose content is in part specified by a frame of reference that's centered on an object separate from the rest of our body and (b) have that content coordinated with the rest of our egocentric experiences without introducing any spatial disunity. Hence, the mere fact that Borgy's experience is structured by three spatially separate origins does not establish that the experience is spatially disunified. To endorse a disunity judgment on the basis of Borgy's having different spatially separated bodies one would have to accept that in the rubber hand and blind torso thought experiments one would experience spatial disunity. And this doesn't seem right.

2.3.3 No Single Frame

In this section, I argue against the view that Borgy's experience is spatially unified because of the absence of a single frame of reference structuring all his experiences. The reasoning behind this way of attacking the spatial unity thesis is this: if Borgy's experiences cannot all be related to a single egocentric frame of reference, then Borgy's experience is spatially disunified. And his experiences cannot be so related. So, Borgy's experience is spatially disunified. But we all agree that it is phenomenally unified. Hence, the spatiality thesis is false.

What is it for any two experiences to be related to a single frame of reference? Consider Figure 2.1 below, composed of six images and two vertical lines:




Figure 2.1 This array could be seen as a single scene, with the cloud-shape in the leftmost position, the sun-shape in the rightmost position, and the remaining shapes: a smiley face, a vertical line, a moon, a rectangle, another vertical line and a heart shape in between. But Borgy could see each pair of images separated by lines by means of a different body, without relating them all to a single frame of reference.

It seems that the above images fall within a single frame of reference. That is to say, there's *a way of specifying the spatial relations of all the images above to a certain object that counts as the origin* (there are actually quite a few such ways). Each shape above, for example, can serve as such an origin.

I call a frame of reference in terms of which one can specify spatial relations between all pictures in a given figure and the origin, *a unifying frame*. If the unifying frame is

egocentric, then the viewer of the figure is the origin of the frame. Now, imagine that Borgy's bodies have the following experiences: body1 experiences the "face" and the "cloud" *only*, body2 experiences the rectangle only and the "moon" only, and body3 experiences the "heart" and the "sun" only. The idea driving Watzl's and Roelofs's interpretation of Bayne's thought experiment could be that Borgy, while still retaining phenomenally unified consciousness of all those images, need not have his experience structured in any way by one unifying frame of reference in terms of which the six shapes above can be related to one another. One way to make this concrete would be to say that the moon could be experienced as occupying some location L1, and the cloud as occupying some other location L2, without the moon being experienced as in any way related to the cloud (e.g. as being to its right). That would be because there needn't be a unifying egocentric frame of reference that structures Borgy's experience, one in terms of which all locations could be experienced. Contrast that with a single-bodied experience one could have looking at Figure 2.2:




Figure 2.2 The images from left to right: a moon-shape, a rectangle, a sun-shape and a cloud-shape

Here, both the "cloud's" location and the "moon's" location are represented in a single experience, and they're also experienced as on the opposite sides of the scene. This is in part because there's a frame of reference that groups them all together, i.e. one that connects the "cloud's" being on the right with the "moon's" being on the left and that frame of

reference structures the visual experience of the scene. E.g., when you visually experience it, it seems to you that *from your point of view*, the “cloud” is to the right of the “moon.” The frame of reference that structures your experience is the unifying egocentric frame. Borgy, in contrast, has no such point of view.

I think it’s clear that Borgy doesn’t have this kind of integration. Borgy sees the face, and he sees the heart, but he needn’t see the face as to the left of the heart. He need not have his experience structured by an egocentric frame of reference with respect to which the spatial properties of both the head and the heart are simultaneously specified.

Hence, Borgy’s experience cannot be structured by a single egocentric frame of reference. It doesn’t follow, however, that this experience would be spatially disunified.

One tempting way to show this would be to appeal to the theoretical distinction between vision for action and vision for perception developed by Goodale and Milner (Goodale & Milner, 1992). According to the distinction, vision for action and vision for perception are supported by anatomically distinct structures in the brain, play different functional roles and track different properties of the environment. Crucially for my discussion here, while vision for action makes use of *egocentric* properties of objects for precise guidance of actions, vision for perception tracks *allocentric*—that is to say viewer-independent—properties of objects.

There is another distinction between vision for action and vision for perception. According to Goodale and Milner’s proposal, it is the latter that correlates with conscious visual awareness. The former carries out its operations unconsciously. This explains shape- and color-constancy phenomena in conscious perception: the objects’ shapes and colors do not seem to us to depend on our perspective with respect to them. Now, however, it seems

to follow from this that the objects' egocentric properties are *not* represented in conscious visual experience. This may strike one as counterintuitive. The content of my experience, after all, seems to be pervaded with how things seem to be from here (recall Tye's description quoted at the beginning of this chapter).

There are ways, however, of capturing this perspectival aspect of the phenomenal character of conscious visual experience without relying on egocentric frames of reference. Berit Brogaard mentions one such way in her discussion of the two visual systems hypothesis:

The information about an object's being over there or to the right of me which forms part of the content of visual perception is arguably allocentric in that it determines, for example, relative distances between an object and the perceiver. It importantly involves the perceiver in determining the axes of the scene, but it does not represent objects in egocentric space. As I perceive the scene in front of me, I perceive parts of my body, for example, my legs and my hands and fingers typing, and I can determine the relation of the clock over there or the soda can to the right relative to my hands and fingers. But my hands and fingers are landmarks in the scene. So, awareness of the position of the clock relative to my hands and fingers is awareness of the allocentric properties of an object (2012).

This suggests that the distinction between vision for action and vision for perception provides an empirically grounded support for blocking the inference from the experiences being structured by no single egocentric frame to them being spatially disunified. This is because the experience Brogaard describes is not structured by any egocentric frame of reference (if you think that the subject represents the hands and legs as her own, and hence views the scene in egocentric terms, we may imagine that, as a result of some psychological condition, the subject is unaware that the hands and legs are hers), and yet it is spatially unified. But this still leaves the proponent of the No Single Frame criterion room for maneuver.

Goodale and Milner's hypothesis does show that it is illegitimate to infer spatial disunity from no unifying *egocentric* frame of reference. But their hypothesis allows for there to be a unifying allocentric frame structuring perception. However, doubt may be cast even on the view that any kind of unifying frame of reference is necessary for spatial unity.

E.g., Jerome Dokic and Elisabeth Pacherie (2006) hold that there could be perceptual experiences that do not involve *any* frame of reference (that is to say, according to Dokic and Pacherie, that these experiences are "perspective-free"):

[I]t is not clear to us that the notion of a frame of reference is needed to explain the relevant phenomena [occurring when you're looking at a picture of a bottle and a chair]. There is another way of seeing the array [in addition to representing the bottle as in front of the chair], namely by representing the bottle as being next to the chair. Of course there is a difference between seeing the bottle as being *in front of* the chair and merely seeing it as being *next to* the chair. The latter relation (x is next to y) is perspective-free and requires no analysis of the objects in component parts (it is enough that they be seen as two different wholes), whereas the former relation (x is in front of y) is perspective-bound and requires such an analysis (2006).

If Dokic and Pacherie are correct, then there could be experiences that are not structured by any frame of reference; there could be "perspective-free" perceptions.

Nonetheless, the bottle- and the chair-representing experiences are clearly spatially unified (the two objects are experienced as occupying the same space if anything is). But this unity is achieved without the two experiences' being structured by any unifying frame of reference. If this is right, then it's possible to have spatially unified experiences without a unifying frame structuring them.

Dokic and Pacherie's claim that *no frames of reference* are required to make sense of the visual experience of the bottle as being next to the chair is more radical than what the two visual systems hypothesis licenses. It is also not supported by empirical evidence.

However, if their claim expresses a coherent possibility, then the No Single Frame motivation for the disunity judgment is flawed.

But could there be an experience of two objects as next to each other, without that experience, *eo ipso*, representing the two objects as being related in some more determinate way? Introspection seems to suggest that the answer to this question is negative. However, while we perhaps can't *visualize* such experiences, we ordinarily undergo them in audition. Imagine hearing two people talk, without being able to see them (they're in another room). The way you experience them is as being next to each other, without representing either as to the left or right of the other. So there are possible spatial experiences where objects are represented as next to each other, without that experience representing them as being related in some more determinate way. Dokic and Pacherie are right, at least when we expand our investigations beyond visual experience.

2.3.4 No Spatial Relations

Another way to motivate the disunity judgment is to say that Borgy perceives *no spatial relations of any kind* between an object O1, seen by body1 only and some other object O2, seen by body2 only. This is a fairly relaxed criterion, as it also allows all kinds of ordinary experience to count as spatially unified. The argumentative step seems obvious: if Borgy does not experience O1 and O2 as spatially related, then his experience is spatially disunified (let's call this the No Relation—No Unity principle, NRNU for short). And Borgy experiences no such relations. So his experience is spatially disunified. But we all agree that it is phenomenally unified. So it's possible to have a spatially disunified and phenomenally unified experience. The spatiality thesis is false.

Bayne clearly states that Borgy perceives no spatial relationships between the objects consciously represented by different bodies. But I don't think this dooms the spatiality thesis. The reasoning above exploits an ambiguity in the term "spatial relation." On one reading, Borgy doesn't experience objects as spatially related, but it doesn't follow that his experience is spatially disunified. On another reading, it's incorrect to say that Borgy doesn't experience objects represented by different bodies as spatially related. But it will take some stage-setting to show all this. The next section is devoted to that.

2.4 Spatial Disunity and Phenomenal Unity are Incompatible

Our spatial experience isn't always uniformly good at representing how things in the world are related to each other, and to us. Sometimes we experience the space around us with great acuity, other times we only get a crude representation of the spatial layout. To see this, compare how you experience the spatial world in vision with how you experience it in audition. It seems uncontroversial that the latter way of experiencing space is impoverished with respect to the former. Typically, vision is *better* for space perception than audition is.

We may say that what's characteristic of a better spatial experience consists, at least in part, in experiencing objects as determinately related (e.g. when I see the cup as to the left of the computer, at such-and-such a distance, and within reach), while in a spatially impoverished experience we experience them as *determinably related* as well (e.g., when I hear two people talking outside my office, without seeing them, I don't represent who's to the left of whom; I do represent them, however, as occupying a dimension that permits this sort of ordering without experiencing *how* they are actually ordered; on the other hand, I do represent determinately the relation of them being next to one another).

To represent two objects as determinately spatially related, relative to a spatial relation type *T* (e.g. being aligned on the left-right axis), is to represent them as exhibiting a particular ordering specified by *T* (e.g. to represent a glass as to the right of the laptop). To represent them as determinably spatially related relative to *T* is to represent them as occupying a dimension in which an ordering along *T* occurs, without experiencing them as ordered in any particular way (as in the hearing example above, I represent the two voices as determinably related along the left-right axis). Less formally, it is to represent *that* they're related without representing the relation in any particular way.

Now, one way to think about this kind of perceptual hierarchy is to adopt an idea from Sean Kelly's Maurice Merleau-Ponty-inspired philosophy of perception. The basic thought that I wish to appropriate here is that our perceptual engagement with the world can be captured by adhering more or less closely to a norm (I remain neutral as to whether this creates some sort of a normative pull that is constitutive of experiencing things; all I'm interested in is that there is a norm at all). Kelly holds that perception necessarily involves awareness of how the current experience could be improved, as well as a compulsion to actually improve it. He summarizes his position thus:

[P]erceiving an object involves ... the perceiver's already being driven to get a better grip on it. When I see an object, it is not up to me to decide to get a better view of the thing I'm looking at. Rather, it's part of being experientially directed toward the object in the first place that this repulsion from lack of clarity is given to me (2010)

Let's say that the norm in question can be said to be the norm of full determinacy (the opposite, say, of the "lack of clarity" Kelly mentions; Merleau-Ponty at one point expresses the norm as "the maximum of visibility" (1998). Our experiences may fail to meet the norm in various ways. Think again about the case of hearing two people talking without seeing

them; my hearing alone doesn't determinately represent who's to the right of whom; it only represents them as next to each other; my experience could be improved—i.e. brought closer to meeting the norm of maximum determinacy—by looking at the two people. Were I to look at them, I would experience how they are actually ordered relative to, e.g., the left-right axis.

Thus, if we arrange our experiences by how well they meet the norm, one end of the spectrum would be occupied by a kind of experience that represents objects as determinately ordered relative to *all* spatial relationship types (perhaps our visual experience comes closest to meeting that norm); it's at least theoretically possible that near the other end of the spectrum are experiences that do not represent objects as determinately related to other things in *any way*. Rather, they represent the objects as determinably related to other things. That is to say, in such experiences objects would be represented merely as occupying a dimension that's characterized by certain kinds of orderings among its occupiers, without representing how these particular occupiers are actually ordered.

We can imagine having experiences falling, to varying degrees, short of the full determinacy norm. But could we undergo an experience from the very end of that spectrum? I think we could. I think there could be experiences with *spatial* content that only represent objects (or locations) as determinably related. But in order to see this, we need a detour through a different kind of, arguably spatial, experience.

Casey O'Callaghan (2010) offers examples of experiences that, I think, could be called minimally spatial: “[1] awareness of a uniform sinusoidal tone at a given loudness ... [2] experience [as of] a uniform-grey ganzfeld as a result of the application of halved table-tennis balls over the eyes ... [3] the uniform feeling of warmth on every skin surface.”

O’Callaghan thinks it at least might make sense to describe such experiences as spatial (as of, e.g., the sound being here). Likewise, in her pioneering work on the sense of smell, Clare Batty observes that smelling things has a similar phenomenological profile to the cases O’Callaghan mentions, at least as far as spatial content is concerned:

If we bracket information gained from movement and any other sensory modality, and consider olfactory experience at-a-time, then we see that any locatedness of these properties — other than simply 'here' — goes as well. In this way, olfactory experience distinguishes itself from vision. Even at-a-time, visual experience achieves a robust form of differentiation that allows the presentation of properties at determinate locations in the scene before your eyes. The spatial presentation of olfactory experience at-a-time remains wholly undifferentiated (Batty, 2010)

These are vivid illustrations of the many ways of experiencing the space around us, including the kind of undifferentiated, minimally spatial experience characteristic of what smelling may feel like. It seems right to say that in undergoing an olfactory experience, I am aware of the smell being simply here, in this place; I do not seem to be aware of it as being *nowhere*, but rather as filling the space.

Mohan Matthen describes the phenomenology of olfaction in a similar way: “In the kitchen ... I smell olive oil and garlic. Where are these odours? Simply *here*. They coexist in a place without combining” (Matthen, 2005), and so does David Chalmers: “an olfactory experience represents that a certain smell is present in one’s environment, perhaps in a certain broad location” (Chalmers, 2006). I think it is safe to assume that even though there isn’t much structure to the space as represented in olfactory experience, some minimal spatial content is undeniably present.

To make the idea of a minimally spatial experience more precise and less reliant on examples, we may speculate that in such an experience the only spatial structure to be revealed would be that of an indiscrete space (a space in which no points are topologically

distinguishable). But it still would be an experience with spatial content (an indiscrete space is still a space).

Let us say then that the previous examples, despite lacking any differentiation with respect to the stimulus, are examples of spatial experience. Now, take two experiences of the kind that Robert Van Gulick describes: “Imagine having just the experience of a dim flicker that passes so quickly that you cannot say where it occurred or whether it was of any given color or shape, or the experience of brief, faint sound whose location and tone one cannot discern. Such stripped-down experience seems possible” (2014). I would add to Van Gulick’s point that while it might be impossible to pinpoint the location of the flash or the sound, they would still seem to us to occur somewhere. The experiences seem to be spatial in some sense. This is because Van Gulick’s examples exhibit more structure than the smell and audition examples countenanced above, since the former allow for more differentiation (neither the flash nor the sound fills up the space around one uniformly, even though in neither experience is it determinately delineated where the stimulus is absent and where it is present). Since the smell experience has some spatial content, an experience that allows for more differentiation than a minimally spatial experience should also count as spatial.

Yet, such an experience as Van Gulick describes doesn’t represent the flash, or sound, as spatially related to anything else one also represents in experience. So how can it count as spatial at all? This is where the notion of being experienced as determinably related may come in handy. The idea could be that in experiencing the flash/sound as in space, I experience it (or its location) as determinably related to other things I already experience as in space, without experiencing *how* they are related. This is just what it is to experience

something as merely occurring in space, to experience it as lying at the opposite end from the full determinacy norm.

It seems extremely plausible that the kinds of olfactory and auditory experiences described by Batty and O'Callaghan are the kinds of experiences that creatures lacking conceptual sophistication may have. It also seems plausible that such creatures could experience flashes and sounds that Van Gulick mentions. But is it legitimate to ascribe to such beings an experience as of something being determinably related to other things? I think it is. To see this, we can appeal to the Autonomy Thesis, introduced, rejected (1992), and subsequently embraced (2014) by Christopher Peacocke. The idea behind Autonomy Thesis is that a creature may enjoy states with non-conceptual content without possessing any concepts at all. The content of an experience representing some things as determinably spatially related could then be taken as the kind of content we can ascribe to the creature without ascribing to it the possession of concepts required to specify the content. In addition, we can explain the creature's behavior by appealing to this content. Imagine that our creature hears such a sound and finds it threatening. Given the spatial content of such an experience, we may expect the creature to start looking around in fear, trying to pinpoint the location from which the stimulus arrived (so that it can, e.g., determine the direction in which to escape). When asked why the creature is scanning the space around it, we can answer that it heard a threatening sound as being somewhere in the space it is experiencing, but did not perceive the sound's location in any precise way. In short, it experienced the sound's location as only determinably related to other locations it also represents. It seems to me that even simple creatures could have such an experience, characterized with the help

of the notion of autonomous content. We even serve some explanatory interests by ascribing such contents to make sense of the creature's spatial actions.

I will now argue that even two such minimally spatial experiences are spatially unified when they are phenomenally unified. This will, finally, allow us to see how to dissolve the ambiguity in the argument for disunity that I'm disputing here.

Let's repeat: the least that a minimal spatial experience entails is that the objects or places perceived are perceived as determinably spatially related to other entities. This is what is meant by the objects' being perceived merely as occupying space. I think that the case of hearing two sounds without representing either as standing in a determinate spatial relation to the other illustrates what I mean well. Thus to ease exposition I focus on sounds in the argument below, but this is not mandatory.

- (1) Part of what it's like to hear one sound is to hear it as coming from space.
- (2) Part of what it's like to hear the other sound is to hear it as coming from space.
- (3) And there's something it's like to hear both (they're phenomenally unified).
- (4) So they're both individually experienced as coming from space (from 1-3).
- (5) If they're both individually experienced as coming from space, then if they are experienced by the subject at the same time (i.e., if they are phenomenally unified), they are experienced as occurring in the same space.
- (6) So, if the two sounds are phenomenally unified, then they are spatially unified (from 4 and 5).

What remains is to justify premise (5). My justification for (5) is that to experience some object *x* as in space is, at least, to experience it as determinably spatially related to other things. Other things include, at least, currently experienced spatial objects (objects

experienced to be in space). Since both x and y are experienced as occurring in space, they are experienced as at least determinably spatially related to other things, *including one another*. And if they are so experienced, then they are experienced as occupying the same space. Notice, though, that although I need not experience *any* determinate spatial relations between the two objects in question, I would still experience them as spatially unified. And this would be in virtue of experiencing them as determinably, rather than determinately related.

The argument generalizes, since the following principle seems almost trivially true: any experience as of an object in space is either an experience as of an object standing in determinate spatial relations to other things or as standing in determinable relations to other things. If so, then either the objects are experienced as occupying the same space in virtue of being experienced as determinately related, or they are experienced as occupying the same space in virtue of being experienced as determinably related. Hence, any phenomenally unified spatial experiences are also spatially unified.

Now we can finally see an outline of a reply to the argument that Borgy's experience is spatially disunified. It is true that Borgy does not represent the objects as determinately related. But it doesn't follow that his experience is spatially disunified. NRNU principle only works when we take "spatial relations" to mean "determinable spatial relations." But it is impossible for Borgy to experience two spatial occupiers without at least experiencing them as determinably spatially related. Thus, the argument above fails.

Let's now see how this would play out in Borgy: he perceives one duck as standing in some determinate spatial relations to his body₁ and he perceives another duck as standing in some determinate spatial relations to body₂. But, it seems to me, he also experiences the

ducks as being in space, as determinably related to other objects. The other objects, I take it, include those that are currently experienced to be in space. And these include both ducks. So both ducks are experienced as standing in determinable spatial relations to each other. So they are experienced as belonging to the same space. It thus appears like Borgy's experience is not a counterexample to the spatiality thesis, *pace* Watzl and Roelofs.

There's a worry that by making spatial unity so much less restrictive, I've made it uninteresting. In particular, the kind of minimal spatial unity involved in experiencing some two things as relatable, but not related, is insignificant compared to a more robust sense of unity that we typically enjoy, and that it's the latter that should really matter in the investigations of the sort I'm engaged in here. But then, one could wonder what principled reason there could be for making this distinction. It might be held that spatial unity is significant only to the extent that ascribing spatial content that represents the objects as being in the same space has some practical import when it comes to explaining actions. But, as I've argued above, even the experiences that Van Gulick mentions may have that kind of import. So this way of drawing the distinction between trivial and interesting kinds of spatial experience cannot rule out the minimal spatial unity I'm talking about.

Moreover, it seems that if I'm right, then a certain way of thinking about the connection between spatial experience, the degree of its integration, and the unity of consciousness stands in need of revision. This would make the minimal spatial unity significant, at least for the view I'll mention presently. The way of thinking I have in mind is pursued by Dainton, in the course of arguing that spatial unity is not necessary for phenomenal unity. Dainton writes:

experiences can be clearly and unambiguously co-conscious ... even when their phenomeno-spatial relations are very indistinct, e.g., the distance and direction of a pair of simultaneously experienced sounds can be very unclear without impairing their co-consciousness in the slightest. If there could be ... beings whose experience is only weakly integrated spatially, why couldn't there be ... beings whose experiences ... unfold in several spatially distinct ... fields? I can see nothing in the character of our ... experience which suggests such beings are impossible. Consequently, even if human beings are incapable of enjoying a spatially disunified consciousness, due to some contingent features of our experience-generating equipment, there is no reason to suppose other types of being need be similarly restricted. (2003)

If what I say in this paper is right, then there is reason to think that we cannot argue from the possibility of experience being less and less spatially integrated to the absence of spatial connectedness. This is because as we gradually decrease the degree to which experience is spatially integrated, the limiting case will not be “several spatially distinct fields” but rather experiences from the very end of the full determinacy norm. However, even when we remove all determinacy from the representation of spatial relations, we will still get an experience of a single space. This is because, as I have argued, even the experience of a spatial relation between two locations that's completely indeterminate remains a spatially unified experience since the locations are experienced as being in the same space. Hence, it seems that Dainton's reasoning runs afoul of the considerations specifically to do with minimal spatial unity. This makes the latter significant, at least for Dainton's view.

2.5 Summary

I have argued that Bayne's thought experiment does not have the significance for the spatiality thesis that Watzl and Roelofs claim for it. I did this by considering, and rejecting, several ways of motivating the judgment that Borgy's experience is spatially disunified. As a consequence, the spatiality thesis remains a viable option, at least in the face of the

thought experiments based on multiple embodiment. In the next chapter I argue that non-perceptual experiences of space (e.g. spatial imagination and spatial memory) are spatially unified with perceptual and somatic experiences. I also argue that there aren't any perceptual experiences that are also not spatial.

Chapter 3: Imagination and Memory Experiences are Spatially Unified with Perceptual and Somatic Experiences

3.1 Non-Perceptual Spatial Experiences

In this chapter I argue that spatial imagination and spatial memory are spatially unified with perceptual experiences. That is, they have the kind of content that *experientially anchors* the imagined/remembered scene in the same space the remembering/imagining subject is also perceiving. This is not achieved via inference through background (maybe geographical) knowledge. Phenomenology of spatial imagination and spatial memory can be explicated in such a way that spatial unity between imagination and memory experiences is achieved without the mediation of background beliefs. Further, I argue that, despite some arguments to the contrary, all perceptual experiences of spatial objects (objects that have some spatial properties, like shape) are experiences as of those objects occupying space.

The kind of imagination I consider is sensory imagination (visualizing, auditory imagination, tactile imagination, etc.); I set aside imagining as supposing a proposition to be true (i.e. propositional imagination).

3.2 Imagining, Remembering, and Perceiving Space

I will talk about imagination first. To focus our ideas, consider the following situation: when on a plane, I visually imagine a friend, Luke, having coffee at a café we both like to patronize. While I do the imagining, I'm also perceptually experiencing the inside of the plane, the taste of Diet Coke, the hum of the engines etc. How are the perceptual experiences related to the imaginative experience?

Well, obviously they are all *my* experiences. They are all parts of my overall unified conscious state. They are phenomenally unified with each other. Typically, I may introspect each of them, and I'm aware of them as my own. But what about the spatial properties and relations represented in perception and the spatial properties and relations represented in imagination? Is there a sense in which the locations represented, respectively, by imagination and perception (the café, the inside of the plane) are experienced as spatially related?

One answer is that they are not experienced as related at all. The phenomenological distinction between my imagination and my perception is partially manifested in that they are experienced as of distinct, different, unrelated spatial environments (true, I may believe that I'm at such and such a distance from the café; but that does not, or need not, affect my overall, spatially disunified perceptual-imaginative experience; if you like, you can change the example to one where I'm imagining Frodo talking to Tom Bombadil; this is imagining a fictional space that I know to be unrelated to any perceivable place in the universe). If things imagined and things perceived are consciously represented as spatially unrelated, then imaginative experiences are not spatially unified with perceptual experiences.

Below, I give three arguments why the antecedent in the above conditional is false. Things imagined and things perceived *are* experienced as spatially related. The first argument is from the possibility of superimposition. The second is from localizability of imagery. The third is from perspectivalness of imagery. The first and the third arguments are the strongest. The third can also be applied to memory experiences.

3.2.1 Superimposing Images

Let us call the space consciously represented in perception “p-space,” and the space consciously represented in imagination, “i-space.” Typically, p-space fills the entire visual field of a subject and even extends beyond it. In contrast, i-space is typically³ confined to the nearest vicinity of the imagined objects.

The first argument why i-space is experienced as connected to p-space comes from the possibility of superimposition. It is possible for you to imagine any object superimposed on the scene you’re currently perceiving, that is, to imagine an object as being in your actual environment (imagining that a leprechaun is jumping on my desk, e.g.). But this merely means that, in imagining something in this way, you allow the perceived scene to impose some constraints upon the imagined object. In particular, you allow the perceived objects to somehow constrain how you imagine the imagined object.

Let’s contrast superimposing imagination with what I’ll call “unconstrained” imagination. Suppose you’re imagining an elephant in an unconstrained way. This means that you don’t let the objects you’re seeing influence in any way how you imagine the elephant. The elephant doesn’t occlude, nor does it get occluded by the objects you perceive. The elephant doesn’t have to bend its knees in order to fit under the tree branches. It doesn’t need anything you perceive as a surface to support its weight. It doesn’t have to get out of other people’s and animals’ way when they walk around the perceived scene. The elephant is sheltered from any interactions with the objects you’re perceiving.

³ This isn’t always the case, e.g. when you imagine being surrounded by snow everywhere, or being in an open sea on a dark and stormy night.

But you can also imagine the elephant as being *in* your perceived environment. That is, as, e.g. floating in the air, or standing on the ground, or dancing on the table, or resting against a tree, or as running away from people you see, or as chasing them. That is, you let the properties of perceived objects (their solidity, opacity, location etc.) to impose *constraints* on how you imagine the elephant.

I contend that, for any imagined object, you can imagine *that object* either as subject to constraints imposed by objects in p-space or as not subject to those constraints. In the former case, you'd experience i-space and p-space as related, and hence as unified.⁴ What about the latter case?

Consider the passage from unconstrained imagination to superimposing imagination in more detail. Suppose you're sitting at a table on a patio, looking ahead. There are, let's suppose, several other tables you also see: some are to the left, others straight ahead, and some others to the right. You also see the ground, some trees, maybe little bit of the sky. Now, as you're sitting and enjoying the view, let's suppose you also visualize a squirrel (in the unconstrained way; that is, you don't recruit your perceptual experience to inform the image of the squirrel in any way); and then, after a while, you begin to imagine the squirrel in the superimposing way; i.e. you imagine it constrained by some properties of some objects of perception; you imagine it as within your environment. Let's say, you imagine the squirrel as climbing the tree you're seeing right in front of you. How do you do it? How do you insert the imagined squirrel into what you're seeing? Well, you might need to do something to the imagined object to make the transition from the unconstrained imagination to the superimposing imagination. You might need to change certain aspects

⁴ Again, you need not thereby *believe* them to be unified or related.

of the imagined object. For instance, part of what you need to do to superimpose the imagined squirrel upon the tree could be to change some aspect, A, of the imagination-experience so that, once you change A, you imagine the squirrel as close to the tree. But perhaps you don't have to change A. Perhaps you'd already imagined the squirrel, in the unconstrained way, such that you don't *now* have to change A in order to imagine the squirrel as close to the tree. So either you have to change A or you don't.

How can one determine whether you need to change A? Whether A needs to be changed depends, it seems to me, on *where* the squirrel, as imagined in the unconstrained way, was with respect to the perceived tree. But since you have imagined the squirrel, in the unconstrained way, as allowing for spatial comparisons with the objects of perception, you must have imagined it as in some as occupying some common overarching framework with the tree, one in which their respective locations can be compared. Since the switch from unconstrained to superimposing imagination is possible for all objects imagined, all objects imagined have to be experienced as sharing some framework with perceived objects. Since this framework allows for comparison between locations, it is some sort of spatial framework. So, the imagined and the perceived objects are seen as sharing some common spatial framework.

Here, I think, two premises are particularly susceptible to objections. First, the universal generalization concerning all objects and scenes that we can visually imagine. Second, the claim that for us to experience something as potentially subject to constraints imposed by objects in p-space, that thing must be already experienced as spatially relatable to places and objects in p-space. I can think of no counterexample to the first premise, however. As for the second, it seems to me that the following holds. If we didn't experience

imagined objects as spatially relatable to perceived scenes, it would seem to require a separate act of imagination to switch between unconstrained imagination and superimposing imagination. But, in fact, no such effort is required. We can have *the same* imagined object switch from one to other kind of imagination in a single act of imagining.

3.2.2 Localizing Imagery

The second argument is the argument from the localizability of the visual image. Think first about a photograph. It represents some scene that has a location in the world. But the photograph itself, taken as a physical object, also has a location in the world, different than the location it represents. Now, for any experience involving visual imagery, we can perhaps draw a similar distinction, between the represented location and the location of representation. There are two reasons to think that the location of representation, when that representation is a visual image, is (a) consciously experienced (b) as a location in p-space. The first is phenomenological: it seems that, when you're entertaining a visual image, you are aware of its location somewhere in front of your eyes, say, or "in your head" (see e.g. (Smart, 1959)). Another reason comes from research on grapheme-color synesthesia (Dixon, Smilek, & Merikle, 2004). The difference in the results of the Stroop test between two kinds of grapheme-color synesthetes has been attributed to the difference in where the visual images seem to be located (inside the subject's head or somewhere in external space). Either way, these images are localizable in the space the subject is also perceiving (we might think of it as egocentric or behavioral space, a la Evans (1985)). Hence, there is spatial unity between i-space and p-space.

One worry about this argument is that it might rely too much on the similarity between visual images and photographs. Perhaps, it might be said, not all visual images need to be

like photographs. It need not always be the case that we're aware of where the image is, with respect to our eyes or head. Again, however, it is not clear how much of a bite such an objection would have. It is difficult to imagine a case where one is entertaining a visual image without being aware of where, in one's visual space, that image is. A different problem has to do with whether the expression "in one's head" should be taken literally to refer to a spatial location. Here, we can point again to the experiment on synesthesia. The difference between the results of the two groups of synesthetes was explained by reference to the fact that those who experienced the image in their heads experienced it as being spatially closer than those who experienced it as in the external world. It was the difference in the distance to the visual image that was taken to explain the difference in task performance speed. This suggests that the "in the head" locution could well be taken literally. And if that's the case, then the location of the image is the location in the same space that the subject is also perceiving.

Another objection is that the argument merely shows the spatial unity between the representational vehicle (the imagined image) used in the imagining and p-space. It doesn't show the latter's connection to i-space. But in accounting for the first kind of unity, the argument above accounts for how every part of the imagined scene or object is experienced as related to objects in p-space. Suppose I'm imagining my dog. In so doing, I'm entertaining a visual image of the dog. Suppose further that I'm experiencing that image a little in front of my eyes. There is a way in which the dog's tail, being on my right, is experienced as closer to the door (also on my right, but in p-space) than the dog's head. Such relations can be experienced for every part of the image I'm entertaining. It is hard for me to see what more could be required from two experiences being spatially unified.

This argument might be seen to rely too much on the photograph metaphor. It might also beg questions in the imagery debate that are too controversial to be begged. Lastly, it might be seen as phenomenologically suspect. Hence, in the absence of some definitive answers in the imagery debate, the argument above is not very effective.

3.2.3 Perspectivalty

The last argument is the argument from perspectivalty of imagery. Let's say that two locations are spatially connected if there is a path from one to the other. And for you to know there's a connection, you don't need to actually know the particular path. E.g., for Sydney and New York to be spatially connected, all that's needed is that there be a path from Sydney to London, and that there be a path from New York to London. I take it that experienced spatial connectedness is sufficient for spatial unity.

Based on these remarks, we may advance the following line of thought: when you imagine something perceptually, you imagine it in relation to yourself. Imagined Luke is "over there" to the right or left, in front of you. That is, you imagine Luke from a perspective, from a point of view (Peacocke (1985), Gregory (2010), Martin (2002), Husserl (1913/1952): apparently a consensus position at least about visualization). At any rate, Luke is imagined as oriented in some way with respect to you. He has a facing surface and a hidden surface; and for that, he needs to be positioned in some relation to your point of view. And this point of view which you occupy when imagining Luke can always be i-spatially related to the imagined object or scene. That is, it's possible to point, roughly, to a location you'd have to occupy with respect to Luke to see him in roughly the way you're imagining him. That is, the point of view from which you imagine Luke is roughly specifiable, given the content of your visualization.

Bernard Williams (1973) and Michael Martin (2002) express divergent thoughts concerning the location of the point of view from which things are imagined. According to Williams, the point of view from which you imagine things is not located within the imagined world (in i-space, as I would say) at all. Martin claims, however, that this view fails to explain our ability to account for certain differences in the visualization's content. Martin says:

One can visualise a red light to the left and a green light to the right. If you now visualise the reverse—a green light to the left and a red one to the right—how you are visualising is different from the first case. Furthermore, this doesn't just reflect a difference in the episode of visualising, rather the two differ because what is visualised is different in the two cases. In the one case the red light is on the left, the green on the right; in the other the green is on the left, the red on the right. But now we can ask what difference need there be in the imagined scene in order for what has been imagined to be different in the two cases? Note, first, that in a world which contains merely two spots of light, there can be no difference between the two situations. The two situations count as different only where there is a point of view relative to which the one object is to the left and the other to the right, or vice versa. So, if we absent a point of view from the imagined scene, then what appears in visualising to be a difference in the scene imagined, and not just a difference in one's state of mind cannot be so. Contrary to Williams's position, ... there must be a point of view within a visualised scene, at least where the visualising involves perspectival elements and those determine aspects of what is visualised. (2002)

I think that Martin's reasoning is persuasive and will assume it in what follows. When you're visualizing something, your point of view is a point of view in i-space, and the range of its possible locations in i-space is determined by the perspectival content of your imagination experience.

Yet, the very point of view that you occupy is itself in real space: it's also the one from which you *actually perceive* things (you can, of course, imagine how things would look like from a different point of view than the one you're actually occupying; but the very act

of imagining would still be done *from* your actual point of view, unless, of course, you lack perceptual experiences—but this is not the situation we’re interested in here).

When you imagine things, you imagine them as they would be related to you. Let’s suppose, with Dominic Gregory (2013), that a visual point of view consists in an origin and a set of labeled axes defining directions (front, left, right, etc.). I claim that it’s plausible to think that the location of the origin of the point of view from which things are imagined, and the location of the origin of the perceptual point of view overlap.




Figure 3.1 The arrows indicate the directional axes, the grey circles indicate objects of imagination, and black squares indicate objects seen.

This is because as you visually imagine things, you imagine how things would look like from some point of view. In visualizing an object, you visualize it in some direction at some distance from you. To do this, your point of view has to be supplied with not only an origin but also with a set of axes defining directions, so that it makes sense to say that the imagined object is in front of you, to your left or right, etc. How these axes are labeled (i.e. where “directly ahead,” “left” and “right” are) determines where your point of view’s

origin would have to be for you to visually perceive the objects just as you visualize them. That is, they determine the location of your point of view in the imagined world. Now, however, you're also entertaining a visual perspective on the world you're visually perceiving. In so doing, you see some things to the left, others to the right, others directly in front. But, just as in visualizing, in perception "the point of view from which one perceives is marked in one's visual experience through it being the point to which the objects perceived are presented—if one can fix the location of those objects, one thereby determine the location of the point of view" (Martin, 2002). And, we should add, one fixes the location of objects by employing the labeled axes.

Now, the "leftishness" of the imagined object, and the "leftishness" of the perceived object are not phenomenologically distinguishable. There is, of course, a phenomenal difference between visualizing something as being on the left and seeing something as being on the left (as Martin points out, for instance, "in imagining something as to the left one does not thereby imagine as in one's actual environment on the left" (2002)). But the direction in which you see/visualize the object does not contribute to the difference in phenomenology. It's the same direction. In the more technical terms, the axis describing the direction "left" is the same axis for both the experience of imagining and the experience of seeing. This is true for all the axes that determine the location of your point of view in the imagined world. They overlap with the axes that determine the location of your perceptual point of view, in the real world. But since they do that, they determine the same location in both cases.

I may, of course, imagine myself to occupy a different spatial position and imagine how things would look like from there. But *how things would look like from there* would

be presented in terms of the axes (some things would be in front, others to the left, others still to the right) that overlap with the axes that determine the location of my perceptual point of view (where some things are in front, to the left, and to the right). The things would be presented to my actual point of view.

So, we can trace a path from an imagined object to the location of your actual point of view (the point of view from which you're doing the imagining which coincides with your actual point of view, and not the point of view you imagine yourself to be occupying, which doesn't coincide in this manner with your perceptual point of view). Then we can trace a path from your actual point of view to a perceived object. Hence, the perceived scene and the imagined scene are spatially connected without being spatially related (or at least, they're not experienced as related in the way that the cup and the book you're perceiving are experienced as related). This is what is meant by saying that the subject's point of view anchors the imagined space to the real (or at least the perceived) space.

Such connectivity with respect to other relations in terms of which unity of consciousness is explained is usually thought sufficient to account for phenomenal unity (e.g. Watzl's attentional connectedness (2014), Masrour's "Leibnizian" connectedness (2014)) I don't see why it shouldn't be sufficient in this case (or, if you don't think that connectedness is sufficient, then this is a weakness my view shares with other plausible views on phenomenal unity; I'm OK with being in the same boat as Watzl and Masrour). We might speak here of experiential spatial connectedness, without actual spatial connectedness; indeed, it might be experiential spatial connectedness without the subject's believing in any actual connectedness.

One problem with this view is that some visual experiences need not involve perspectival-ness (cf. (Dokic & Pacherie, 2006) and section 3.3); however, cases such as described by Dokic and Pacherie (where two objects can be visually experienced, in a “perspective-free” manner, as merely being next to each other) can be reinterpreted as involving the representation of two locations only, in a roughly topological way: inside (corresponding to where the subject experiences himself to be) and outside (corresponding to every place that is not where the subject is). Then, any two objects seen as “outside” would be seen as next to one another (i.e., as both occupying a location that is not the subject’s location) without being seen in any other relation to one another and without being seen as having a particular orientation. Perhaps ordinary human perceivers can’t visualize such experiences.

Even so, the inside, in this case, would coincide with the point of view from which typical visual perception occurs. We might say that while there are directional axes in ordinary visual perception, the only kind of spatial relations represented in the kinds of experiences Dokic and Pacherie mention would be topological relations. However, since to represent something to the left or right of me requires representing it as out there, this complication raises no problem for my view.

3.2.4 Memory Experiences

An experience of remembering seems phenomenologically similar to an experience of imagining (cf. Byrne (2010)). One reason for thinking so is that both in memory and in imagination one can adopt different perspectives on the event imagined/remembered. One can imagine/remember from 1st-person point of view; and one can do so from a 3rd-person point of view (see e.g. (Rice & Rubin, 2011)). Using the same ability to switch perspectives,

to change the way from which you're viewing the scenes seems to require some common component in both memory and imagination experiences. At the very least, it requires that the events be imagined/remembered from *some* perspective. Given this is right, then at least the perspectival argument can be repeated for the recollective experiences as well as for the imaginative experiences.

As a result, I take it that memory experiences should lend themselves to a similar treatment as imagination experiences. Hence, they present no additional problem for spatial unity of experience.

It might also be held that some memory and imagination experience involve (only) an experience of conscious thinking about the situation you're imagining/remembering. Such experiences don't seem impossible (in fact, it's how it seems to me that I imagine/remember most things). Whether such experiences exhibit spatial phenomenology hangs on whether conscious thought exhibits spatial phenomenology. I will argue in Chapter 4 that they do. If I'm right about thoughts, then such intellectual kinds of imagination and memory experiences don't provide a counterexample to the thesis that phenomenal unity requires spatial unity.

3.2.5 Unity without Distance?

Consider next the case described by Mohan Matthen (2014), seemingly showing spatial disunity in imaginative experience: “[O]ne can visualize two objects without imagining how they are related to one another. For instance, one might image the Eiffel Tower and the Tokyo Tower side by side in order to compare them, with the Eiffel Tower on the left. Yet, one might not imagine a specific distance between the two landmarks.” It's unclear to me whether Matthen intends this observation to count against the spatial unity between the

two experiences (of the Eiffel Tower and the Tokyo Tower). After all, one could also trace a path from one tower to another through the subject's point of view. Secondly, the towers could still be imagined as spatially related without imagining a specific distance between them. They could be imagined as merely determinably related, with a determinable, but not determinate distance between them. Lastly, for the example to work, the representation of one tower must be spatially related to the representation of the other tower.

This concludes my discussion of memory experiences and imagination experiences. They don't seem to present a challenge to the spatial unity view.

3.3 Non-spatial Seeing

Another problem for the view I'm proposing arises from considerations of what the experience of some patients with Balint's syndrome is like. One of the best-studied patients in this area is known in the literature as RM. RM's performance in various experiments suggests that his ability to perceive spatial objects does not require the ability to perceive "space as such" (Schwenkler, 2012). In particular, RM seems unable to locate the objects he's perceiving anywhere in objective space, and yet, he's able to accurately perceive the objects' shapes. This has led some theorists to postulate that RM doesn't perceive the objects as in space at all. This view threatens the spatial unity thesis, because if RM doesn't perceive an object as in space, then he cannot perceive it as spatially unified with anything else.

3.3.1 Does RM Fail to Perceive Space?

The argument, proposed by John Schwenkler seems to be premised on the idea that RM fails to perceive objects' locations and orientations, but he does perceive the objects themselves. Schwenkler describes the case thus:

RM had suffered a pair of strokes that resulted in severe damage to his posterior parietal cortex, and as a consequence was unable to localize the objects he saw. For example, when shown a display with a target at one of five locations along the vertical or horizontal meridians and told to report whether the target's position was up, down, or center (in the vertical blocks) or right, left, or center (in the horizontal ones), RM averaged only 70% correct across all conditions. Similarly, when instructed to judge the relative position – left or right in one block, up or down in the other – of an „X“ with respect to that of an „O“ that was also presented on the screen, RM was correct approximately 50% of the time, a performance no better than chance (Friedman-Hill, Robertson, & Treisman, 1995). RM could detect the target stimuli well enough; he just couldn't tell where they were, either on the screen itself or with respect to other things.

Schwenkler concludes that RM doesn't perceive the objects as arranged in a larger spatial framework at all: “his experience was of a wholly unimaginable sort: not of oriented spaces that were cut off at the boundaries of the things that occupied them, but of things without *spatial locations* or *orientations*. RM perceived shapes that did not appear to be in space at all.” RM's experience thus falsifies the thesis that experience of objects requires a representation of space as some sort of overarching framework. Indirectly, it also provides a counterexample to the spatiality thesis, endorsed by this work. However, the conclusion that Schwenkler draws does not actually follow. RM's results are consistent with the view that he does perceive objects as in space, but that the kind of information about their location he is able to extract is severely limited.

Furthermore, Schwenkler's interpretation of what RM's experience is like is inconsistent with a range of other experimental results and observations made on RM. Firstly, Robertson et al. (1997) report that even though RM did poorly in identifying which of the array of objects was in motion, he nonetheless “frequently reported seeing motion in displays that included no moving objects ... He would complain that the stationary letter whose color or identity he was naming was drifting about on the screen.” Now, it does seem impossible to experience motion of an object without experiencing a change of the

object's position relative to some reference frame. RM couldn't have experienced the letters moving, had he not experienced them as changing locations. It is difficult to see how such an experience could be accounted for while maintaining that RM is unable to see the moving object as in space.

Another set of results that appears inconsistent with Schwenkler's view is the data on reaching in RM (Baylis & Baylis, 2001). In the experiments, RM is asked to reach to one of the twelve locations at which a point of light appears. He is better than chance, though his results still are low for the task that's relatively easy. The results of these experiments suggest that RM's grasp of locations in egocentric space is fairly good. But being able to locate an object in egocentric space seems to imply some ability to experience the object in relation to oneself. And hence, as occupying the same spatial framework.

However, care should be taken when drawing conclusions about spatial perception from the tasks requiring the subject to perform an action. Given the empirical discoveries concerning the two visual systems hypothesis (Goodale & Milner, 1992), the reaching data should not be used as evidence for what kind of experience RM is a subject of. This is because, roughly speaking, the system guiding visual action and that responsible for conscious vision function largely independently of one another, and the results from experiments examining visually-guided action should not be used to form conclusions about visual experience (or, more generally, visual perception).

Interestingly, Baylis & Baylis have also examined RM's ability to *report* on perceived stimuli without engaging in reaching behavior. In one experiment, RM was required to merely name (without reaching) the location at which the stimulus appeared (in egocentric terms) and, again, his errors, though multiple and more frequent than in the reaching task,

are mostly mistakes concerning depth (RM gets the direction—left, right, center—correct much more often than chance, but he frequently seems incapable of judging correctly whether the objects are nearer or farther away from him). Mistakes in direction but not the depth of the stimulus were rarer, mistakes in both, rarer still. There then seems to be a good case for thinking that RM perceives the object at least in relation to himself, and even that, in part, the perception is accurate (after all, he was above chance, though not as good as regular perceivers, in determining whether the object appeared on the left or right). There is, then, at least an implicit representation of one object (himself) spatially related to another (images on the screen). And for this, it seems that representing these objects as occupying a single spatial framework is required.

Similarly, another set of experiments conducted on RM by Phan et al. (2000) suggests, in the experimenters' own words, that RM "has some ability to localize stimuli with respect his own body" and that he is able to use an egocentric frame of reference to do so. This, in turn, suggests that RM can experience objects as in a place that's distinct from the one his body occupies. Plus, the use of an egocentric frame of reference implies some ability to represent objects as occupying the same space as his body does.

It might be objected that the experiments Schwenkler bases his interpretation on were made earlier than the ones conducted by Baylis & Baylis and Phan et al. RM's condition may have improved by the time he was examined the next time. Or, it might be that the nature of the tasks was different enough that the experiences evoked were correspondingly different. These objections are well-taken. It seems to me, however, that one reason that may be offered to offset the strength of these objections is that we may explain all the experimental data (i.e. those appealed to by Schwenkler, and those related above) by

postulating impaired, but not non-existent representation of space. In this way, RM's improvement over time would be one of degree (increasingly more detailed representation of space) rather than in kind (some representation of space vs. no representation of space). We would have to give divergent explanations of the two sets of results, if we work along the lines that Schwenkler recommends. A (more) unitary explanation seems better than a non-unitary explanation. A less radical explanation seems better than a more radical one. Thus, in the next section I propose an interpretation of RM's initial results that allows for seeing his improvement as a gradual progression.

3.3.2 Growing Determinacy of Spatial Vision

I take it that Schwenkler's reasoning proceeds along the following lines: RM consciously represents an object at least well enough to detect it and recognize it. But he denies knowledge of the identified object's location. Nor is he able to accurately detect the object's orientation. RM, then, has a conscious experience of objects, and of their intrinsic spatial relations, without being aware of them as in space at all.

First of all, it seems to me that it doesn't in general follow from one's inability to locate an object with respect to another that one does not perceive the object as in space. For consider the following situation: you have your eyes closed. You hear two people talking. The spatial content of your auditory experience doesn't specify the spatial relation between the two voices. You can't tell who's to the right of whom. That is, you seem to perceive two objects without being able to perceive how they're spatially related to one another, what their orientation is, etc. But it doesn't seem to follow that you don't perceive them as in space at all. At the very least you hear them both as *out there*, even though you can't really tell where they are with respect to one another. You'd score poorly on a test

measuring your capacity to locate one sound with respect to the other. However, you wouldn't experience the sounds as not in space. It is possible that RM's visual experience is similar to ordinary auditory experience, as far as the spatial content it provides is concerned.

The experimental results and the patient's reports Schwenkler appeals to are consistent with the idea that RM's visual experience is in some respects like our regular auditory experience; it may only represent some spatial relations without representing any other spatial relations; in particular, it might represent location broadly enough that little information useful to do well in the experimental paradigm could be extracted from this representation. Schwenkler seems to assume that to represent something visually as in space, we need to represent its precise location. But this isn't a necessary truth. Spatial representation is varied. The confusing nature of RM's reports might reflect the unusual poverty of his visual experience of space, rather than the complete absence of the representation of space as such.

I take it that anyone in RM's situation would be confused by not being able to perceive the precise location of objects one used to be able to localize precisely. This inability would presumably lead some rational persons to make reports that RM is quoted as making ("I can't see where it is"); it would, it seems, make a rational person particularly reluctant to guess the location of the stimulus. Compare RM's predicament with the following situation. If all I know about the location of a bear is that it's outside (that's all I heard on the radio, say), and then you'd ask me where the bear is, I could reasonably say that I don't know, especially if I assumed you meant a specific spatial location, or a specific relation to some other place e.g. our favorite spot in the woods, or even with respect to my own

point of view or your body or whatever. Further, if I kept switching radio stations only to find out from their reports that the bear is outside (perhaps each formulated in a different way, or sometimes in a different language), and you kept asking me for (what I'd assume would be) its precise location or orientation, I'd continue denying knowledge of that. It would not, however, mean that I'm not aware of the bear as in space at all. My representation of the bear's spatial relations would be impoverished but not nonexistent. The data Schwenkler cites (RM's poor results when determining the locations of objects) is consistent with such an impoverished way of representing space as such.

Now I think RM's results that Schwenkler points to could be described in a similar way. Sure, RM's representation of space is severely impaired. But, arguably, it could be that it's impaired in roughly the way my representation of the space surrounding the bear is. RM experiences the objects as occupying determinable locations, without perceiving determinate ones just like I know that the bear is at some determinable location without knowing where it is. The difference is, of course, that I'm learning about the bear through testimony, and RM is learning about the object(s) through sensory experience. But that doesn't seem to show that *what* (the kind of information) we learn has to be different.

In other words, an alternative interpretation of RM's case is that there seems to be representation of objects as in a determinable, but not determined, location in egocentric space, perhaps against some undifferentiated background. This interpretation shows that there is some experience of at least some spatial relations between objects (or locations); and to experience spatial relations between different objects, it seems necessary to represent them as occupying some common encompassing framework. Hence, RM's experience is not a counterexample to the thesis Schwenkler sets out to refute. A fortiori,

it's not a counterexample to the spatiality thesis. It's possible to represent an object as in space without representing its location in a determinate way.

One problem with my interpretation is that it ignores the following aspect of RM's performance: he is able to recognize objects (such as letters or short words); hence, he is able to represent fairly detailed *intrinsic* spatial relations (those holding between parts of the object) without being able to recognize the objects' orientations. This leads Schwenkler to conclude that RM doesn't perceive the objects as in space at all. However, correctly identifying a complex object does not require correctly identifying its orientation even in ordinary experience. Ordinary humans are fairly good at categorizing such objects as animals even in far peripheral vision. E.g. subjects performed way above chance even for stimuli appearing at eccentricity 51.5° when the task was to categorize the target stimulus (Thorpe, Gegenfurtner, Fabre-Thorpe, & Buelhoff, 2001). However, in an experimental paradigm by (Sally & Gurnsey, 2003) the subjects' judgments of the spatial *orientation* of the stimulus was poor (at chance) when the stimulus was removed from foveal vision even by as little as 8° . This suggests that the ability to correctly categorize a complex stimulus and the ability to correctly report stimulus orientation could come apart even in ordinary (peripheral) vision. But it doesn't seem right to say that the stimuli that are correctly categorized are therefore experienced as not in space at all. Hence, RM's results in which he was able to correctly identify objects without being able to correctly determine their orientation don't show that he perceived the complex objects as not in space.

So, there is a competing interpretation of RM's results, roughly along the following lines. In the initial set of experiments RM is able to perceive locations only determinately (perhaps merely as out there). This ability improves with years, as he is able to perceive

the locations as increasingly determinate (in front, to the left or right etc.). This is a difference in degree, rather than the one in kind that accepting Schwenkler's interpretation would force us to postulate. Other things being equal, it seems more parsimonious and hence preferable. Crucially for my overall view, it also allows us to see that RM's experience is not a counterexample to the spatiality thesis. RM's perceptual experience is spatially unified.

3.4 Seeing without Connecting

It has been suggested that when you are experiencing phosphenes you experience them as not belonging to the same space as the objects of perception. That is, even though a phosphene experience has some spatiality, it doesn't appear to the subject that the phosphene is in the same space as the things perceived. In this section I argue that this is a misguided view.

When it comes to experiences such as phosphenes, "seeing stars" or viewing afterimages, we seem to be faced with two intuitive judgments as to the spatial phenomenology of these experiences. On the one hand, it doesn't seem to us as if the phosphenes etc. are external to us in the same way that ordinary objects of perception are. There's a phenomenological difference, as Susanna Siegel (2006) puts it, between seeing stars and "seeing" "stars." On the other hand, it is not as if phosphenes and "stars" are experienced as lacking spatial properties and failing to stand in spatial relations altogether. These experiences do seem to have some spatial relational content with respect to the subject: as Siegel observes, they appear in front rather than in the back, for instance.

One way to get the result that a phosphene experience is spatially connected with other (perceptual) experiences, while respecting both intuitions is, again, through the notion of a

path. Just as the imagined Eiffel Tower is imagined from a point of view, so does a phosphene occupy a rather determinate location with respect to your body (right in front of your eyes). Again, this determinate location is a location with respect to your point of view. The same point of view you're enjoying when perceiving things. So, there is a traceable path from the phosphene, through you, to what you're perceiving.

However, it seems to me that going about spatial unity in the phosphene case by drawing a path through your point of view seems unnecessarily complex to account for the phenomenology. But this should give us pause. Why is that the case? One explanation why the path explanation seems unnecessarily complex here is that phenomenologically, it does seem that the phosphene appears closer to my eye than, say, the tree I'm seeing. My conjoint experiences (perceiving the tree and experiencing the phosphene) provide me with enough resources to make these sorts of spatial comparisons. This seems sufficient to claim that the two experiences exhibit spatial unity.

There is a way, then, of describing the phosphene experience in which the phosphene is experienced as related to objects the subject is currently perceiving. There's a sense in which it seems closer to your eye than, e.g., the tree you're also seeing (even though the phosphene doesn't present itself as three dimensional; that, by itself, should not make us judge the phenomenology of having phosphenes to be, somehow, endowed with some mysterious, extraordinary spatiality; we see rays of light and reflections all the time but we don't suppose they occupy some different space just because they're experienced as two-dimensional; seeing a phosphene is very much like seeing a patch of light).

Perhaps to account for the phenomenology of seeing phosphenes not to occupy the perceived space, one needs to make a distinction between a seeming state and an

experiential state. Perhaps it is due to the seeming (a more “intellectual” part of the overall experience) rather than the sensation that the subject takes the phosphene not to occupy the same space as the trees and rocks she’s also perceiving (if we assume that seeming contents and experience contents could come apart). Or perhaps, she simply judges the phosphene to not be in the same space, despite appearances.

I conclude that neither imagination experiences, memory experiences nor phosphene experiences constitute counterexamples to the spatiality thesis. Given plausible assumptions concerning the kinds of spatial phenomenology that memory, imagination, and phosphene experiences have, it looks like these experiences are phenomeno-spatially connected with perceptual experiences, either through a path via the subject’s point of view, or merely in virtue of appearing, phenomenologically to be in the same space as perceived objects (and this, in turn, in virtue of its location being comparable to that of other objects). Obviously, sophisticated perceivers are aware *that* the phosphene experience does not represent any object in space. This might be in virtue of a judgment that the phosphene is not in objective space. But it doesn’t show that the phosphene doesn’t *appear* to be spatially comparable with other objects.

3.5 Summary

How does the above relate to the principle that every experience as of an object in space is, *eo ipso*, an experience as of an object at least determinably related to other objects also experienced to be in space? Firstly, the correctness of Schwenkler’s interpretation of RM would make it the case that the principle doesn’t apply to all object-perception. I hope to have laid that worry to rest, however. RM seems to see objects as occupying the same space he’s also occupying.

Second, how does the principle connect with the ideas advanced here about the phenomenology of remembering and imagining? I think that the idea of a path is one (but not the only) way in which we may think of two spatially experienced objects being experienced as determinably, but not determinately, connected. Imagined, or remembered, friend, appears at some location that need not be experienced to be determinately related to perceived locations (perhaps this accounts for what some would see as the oddness of the question “Where is your imagined elephant located with respect to the desk that we’re both seeing?”). But, via the path, the imagined and the perceived get to be experienced as determinably related (this is also what accounts for the phenomenal difference between *perceiving* two things as determinably spatially related on the one hand, and experiencing an imagined location as connected to the perceived locations; in the first case we don’t need to go through any paths).

That’s what I mean by saying that the subject’s point of view anchors the imagined space to the perceived space. If I’m right that this is a species of representing spaces as determinably related, then the principle works across all sorts of spatial experiences, not only in perception. It is just that the model for connecting whatever’s represented by a mental image with whatever’s perceived is more complex than the model connecting, e.g., Borgy’s perceptual experiences.

Chapter 4: A Case for Spatial Cognitive Phenomenology

4.1 Do You Experience Your Thoughts as Occurring in Your Head?

So far, I have established that perceptual representations of space are spatially unified with spatial imagination, spatial memory, and other kinds of spatial, non-perceptual experiences. Now is the time to show that conscious states taken to lack spatial phenomenology, e.g. conscious thoughts and moods, are also, in fact, spatially unified with the experiences I have discussed so far. Accordingly, in this chapter I argue for the view that conscious thoughts⁵ exhibit the kind of spatial phenomenology that allows for spatial unity with other spatial experiences. More precisely, I claim that whenever we engage in the activity of conscious thinking, there is an aspect of the phenomenology of conscious thinking that makes it seem to us as if that activity, or its products (i.e. individual thoughts), has (or have) some apparent location. Further, the phenomenology in question need not be dependent on the content of the thought, but rather, it could be dependent exclusively on the kind of propositional attitude (consciously thinking that...) entertained with respect to that content.⁶

⁵ By “thoughts” I mean cognitive mental attitudes with propositional content, such as entertaining, doubting, hoping, supposing etc. I want to distinguish conscious thoughts, for the purposes of this paper, from such mental states as imaginings, rememberings etc.

⁶ This claim is not necessary for my position. For it could well be that the phenomenology is due to the thought’s content anyway. It is just that the content in virtue of which thoughts have spatial phenomenology (e.g. “that this thought is occurring in my head”) is more helpfully thought of as peripheral (because usually less important to the thinker) to the central, core content of the thought (e.g. “that $2+3=5$ ”).

Spatial cognitive phenomenology, then, admits of some variety: it is present not only when one's individual thoughts seem to occur somewhere, but also when one's own activity of thinking (i.e. producing a string of thoughts) seems to one to occur somewhere. Thoughts also have spatial phenomenology, on my view, when their *vehicles* exhibit spatial phenomenology (e.g. when one "thinks in images" and is conscious of these images' location).

It's important to note that my use of the term "vehicle" is closer to that of, e.g., Michael Dummett (1994) and Jose Luis Bermudez (2003) than to that of, e.g., Susan Hurley (1998). For Hurley, the vehicles of mental states "are usually thought of as neurophysiological or computational states or processes located in the brain" (1998). This makes them "subpersonal," that is to say, they are more usefully thought of as states of some subsystem within a person (a brain, or some part thereof), and hence they would not figure in folk-psychological explanations of behavior. Vehicles, on this understanding, are also not something we can customarily become aware of.

The way Dummett and Bermudez use the term "vehicle" appears to differ from the Hurley use, although neither Dummett nor Bermudez is particularly explicit about just what the "vehicle," on their understanding, is.

In Dummett's work, the term occurs in the context of distinguishing real thoughts from proto-thoughts. This distinction is captured partly by appeal to the thoughts' vehicles:

A car driver or canoeist may have rapidly to estimate the speed and direction of oncoming cars or boats and their probable trajectory, consider what avoiding action to take, and so on: it is natural to say that he is engaged in highly concentrated thought. But the *vehicle of such thoughts is certainly not language*: it should be said, I think, to consist in *visual imagination superimposed on the visually perceived scene*. It is not just that these thoughts are not in fact framed in words: it is that they do not have the structure of verbally expressed thoughts. But they

deserve the name of “proto-thoughts” because, while it would be ponderous to speak of truth or falsity in application to them, they are intrinsically connected with the possibility of their being mistaken. (Dummett, 1994; my italics).

Bermudez and Dummett distinguish such proto-thoughts (“vehicled” in spatial imagery) from thoughts “vehicled” in a public language. Dummett says that, in contrast to proto-thoughts that have spatial imagery as their vehicle, *language* is the vehicle of “full-fledged thought.” Similarly, Bermudez speaks of propositional thoughts as being vehicled by sentences of inner speech, and claims, furthermore, that whenever we’re introspectively aware of such thoughts, we’re aware of them as sentences in a public language. He writes: “we are not, I think, ever *conscious* of propositional thoughts that do not have linguistic vehicles. When we are conscious of propositional thoughts we are conscious of imaged sentences” (Bermudez, 2003).

In general then, vehicles, on the reading I’m borrowing from Bermudez and Dummett, *are personal-level states the having of which enables the contents of the subject’s thought to become phenomenally conscious (or, in more colorful terms, to bring the thought’s content before one’s mind)*. Visual imagery is one example of such a vehicle (as in Dummett’s vignette); auditory imagery is another (when you’re aware of having a thought in virtue of engaging in silent speech). I think it is reasonable to ascribe this understanding of the term “vehicle” to Bermudez in particular, because of his two claims: one, that our introspective awareness of propositional thoughts is always awareness of them as sentences in the public language, and two, that public language is a vehicle of propositional thoughts. Hence, my use of the term “vehicle” follows established, though not the most common, usage.

In general, it looks like in some cases the spatial phenomenology of thoughts is more like that of pains and other bodily sensations and not that of visual perception, in the following respect: pains and other bodily sensations have spatial phenomenology in virtue of representing their own apparent location. Perceptions, in general, have spatial phenomenology in virtue of representing the location of objects distinct from perceptions themselves. Other cases of spatial phenomenology involved in conscious thinking, however, might be more difficult to place on the self-location side of this divide. Suppose that conscious thinking involves “thinking in images,” and suppose further that thinking in images is a quasi-perceptual process in which the images appear perspectively. To view an image perspectively is to view it from somewhere. Hence, conscious activity of thinking in images would have some spatial, or locational, phenomenology, perhaps independently of the thought’s content. It might then be held that the quasi-perceptual phenomenology of experiencing the vehicle of the thought’s content (e.g. when seeing images or hearing sentences in one’s head) endows conscious thought with spatial phenomenology. I intend the discussion that follows to cover all these kinds of spatial cognitive phenomenology.

One can distinguish a strong and a weak version of the view that conscious thoughts exhibit spatial phenomenology. On the strong reading (what I’ll call the strong cognitive spatiality thesis), conscious thoughts *always* exhibit spatial phenomenology. That is, there is no instance of consciously thinking that *p* that does not make the subject aware of some spatial location (independent of the thought’s content). On the weaker reading, some

conscious thoughts may be held to have spatial phenomenology while other thoughts don't. I'll call this the moderate view about spatial phenomenology of thought. My goal is to show that the strong cognitive spatiality thesis is more plausible than the moderate view, or the wholesale denial of spatial phenomenology to thoughts.

I admit, however, to having no knockdown argument against the moderate position. My aim, therefore, is first to argue that the strong cognitive spatiality thesis is a more reasonable starting position than the moderate view. If this is right, then the strong cognitive spatiality thesis should be given up only if there are good objections against it. I examine several such objections in the course of this paper, and argue that they don't compel us to give up the thesis as a reasonable default position.

I call the view that postulating spatial cognitive phenomenology is a reasonable default position the dialectical advantage thesis. I offer three I offer three reasons for accepting it: (1) naïve introspection alone does not allow for the resolution of the debate between the strong view and the weak view. Predictably, the proponents of the strong view as well as the proponents of the moderate view will hold that introspecting their own experience supports their own position. But the moderates seem to occupy a more privileged dialectical vantage point here. Their own introspection might serve, after all, as an existence proof of conscious thoughts without spatial phenomenology (thus directly contradicting the strong cognitive spatiality thesis). I will argue, however, that relying on introspection alone, we cannot advance the debate between the proponent and the denier of spatial cognitive phenomenology. At most, we'll reach an impasse.

My second reason why the strong thesis is dialectically privileged is that spatial cognitive phenomenology goes hand in hand with a broader range of approaches advocated in the cognitive phenomenology debate than its denial. The general thrust of the argument why the strong view on spatial cognitive phenomenology enjoys a dialectical advantage in that context is that it is easier to accommodate with a wider range of positions in the more general debate. Cognitive spatial phenomenology goes hand in hand with sensory-based views on cognitive phenomenology and does not seem to be inconsistent with views that postulate proprietary cognitive phenomenology. It is more difficult to see how a denier of spatial phenomenology could accommodate as broad a variety of views in that important debate.

While it is still open to the denier of spatial phenomenology to endorse a view in the broader debate that chimes better with his or her denial, the fact that proponents of spatial cognitive phenomenology have a broader array of views to ally themselves with seems like a reason to assign dialectical advantage to their view, small though it is.

Lastly, spatial cognitive phenomenology enables us to make clear sense of how to describe thought broadcast phenomena (one of the so-called first-rank symptoms of schizophrenia). It is hard to see how the deniers of spatial cognitive phenomenology (of either moderate or extreme variety) can handle some cases of thought broadcast in a principled way. If this is right, then spatial cognitive phenomenology has a clear dialectical advantage over its denial.

I hold that if one accepts (1)—(3), then one ought to adopt the view that conscious thoughts always exhibit spatial phenomenology as default, absent compelling objections. Therefore, if one intends to hold either the moderate view or if one wants to deny spatial

phenomenology to thoughts entirely (the positions aren't always clearly distinguished), one should offer objections against the strong view. Otherwise, one is not entitled simply to assert that conscious thoughts lack spatial phenomenology, either sometimes or always.

4.2 Spatial Phenomenology in Conscious Thought

Consciously entertaining thoughts is a phenomenally conscious process. If consciously entertaining thoughts involves spatial phenomenology, then among the phenomenal properties instantiated when the subject is consciously entertaining thoughts, there is some phenomenal property the instantiation of which makes the subject phenomenally conscious of some location, and, moreover, the experienced location need not depend on the content of the thought in question.

To get clear about the kind of spatiality I have in mind when I talk about spatial phenomenology of conscious thoughts, I will consider here two related objections to the view that thoughts do have spatial phenomenology. We might call the first the Humean objection. The second objection builds upon the Humean one. Our Humean objector may argue as follows: if thought really exhibited spatial phenomenology, then it would make sense to speak of sensing one of one's conscious thoughts to be to the left or right of another. But it makes no sense to speak in this way. It's absurd to think that you could ever go through, e.g., a conscious piece of reasoning where the consciously entertained premises seem to you to be to the left of your consciously entertained conclusion. And even if it did sometimes make sense to think that this could happen, (in some cases of pathological spatialization of experience – see Parnas et al. (2005), it seems wildly implausible to suppose that our ordinary cognitive phenomenology involves such properties. So, conscious thought does not exhibit spatial phenomenology.

Another objection along similar Humean lines could be that if thoughts are experienced as occurring in the subject's head, then it should make sense to speak of experiencing one's thought as moving around in the head, or at least as capable of moving around. But it makes no sense to speak in this way.⁷ So, thoughts are not experienced as in the subject's head.

Where I think these two objections falter is in assuming that the apparent location of conscious thoughts conveyed by spatial cognitive phenomenology is fairly precisely specified. On reflection, this approach to spatial phenomenology might be seen as too restrictive. We should be more broad-minded about how our inner experiences' location could be specified. My proposal is that we understand location phenomenology of thought to admit of being less precise than the kind of spatial experience where such relations as "is to the left of" or "is behind" are experienced. In particular, I think it might be worth conceiving of some of the spatial properties involved in spatial cognitive phenomenology as *topological* properties.

In the next section I explain what I mean by topological properties and show how appealing to them helps address the two objections. But the general thought is that the consciously experienced spatial properties that contribute to what it's like to have conscious thoughts are more coarse-grained than such spatial properties as "being to the left" or "being behind." Rather, the consciously experienced spatial properties could be exhausted by, e.g. the topological property of being outside or inside.

4.2.1 Topological Properties

The topological property that would be of most interest to us is the property of enclosure. To see what the property is, imagine an ordinary triangle. The triangle has a

⁷ This objection can be resisted even at an earlier stage. After all, it does make sense to say that thoughts are running through my head.

number of properties. It has three internal angles, sides of a certain length and so on. One of the triangle's properties is also that it divides the plane into what is inside the triangle and what is outside it. Now, topology deals with those properties of figures that remain constant under bi-continuous one-one transformation⁸ (cf. e.g. (Flegg, 1974)). The property of dividing the plane into an inside and an outside would be preserved under such a transformation, and hence counts as a topological property. The size of the angles or the length of the sides would not. Why? Intuitively, bi-continuous one-one transformation is the transformation that preserves the properties of figures when they are bent, compressed, stretched or twisted in any way that ensures that the sides are not "cut" (e.g., a triangle could be transformed into, and hence is topologically equivalent to, a circle). What has to remain the same, in the kinds of transformations allowed by topology, is that the points that are "near" before the transformation remain "near" after the transformation. That is, the points in the first figure are mapped onto points in the transformed figure in such a way that those that are inside remain inside, and those that are outside remain outside. Given that distance is not preserved under transformations allowed by topology, it is more useful to think of nearness in terms of remaining inside, or outside, the plane figure.

The relationships of distance or shape are not preserved under transformations allowed by topology. Similarly, the precise location of a point, given by, say, Cartesian coordinates, becomes irrelevant when topological properties are concerned.

⁸ E.J. Green offers a very helpful explanation of topological properties: "A topological property is any property that is preserved under all topological (i.e., one-to-one, continuous) transformations. Topological transformations are often called "rubber sheet" transformations, because they include all the deformations one can apply to a rubber sheet—e.g., twisting, stretching, bending, etc. However, they do not include tearing an object in two, poking holes in an object, "filling in" the holes of an object, or "gluing" pieces of the object together. Topological properties include connectedness, an object's number of holes, and an object's property of being inside or outside another object. Because of the generality of topological transformations, any two solid figures—e.g., a ball and a block—are topologically equivalent" (Green, forthcoming).

Now, I claim that the spatial phenomenology involved in the having of conscious thoughts is in some cases more adequately captured when we take it to involve topological properties rather than, say, metric properties. So, it might be that (some) inner experiences have the spatial phenomenology to the extent that they are felt to occur inside (the subject, say) without, at the same time, being felt to occur in a particular location within the subject (at a certain distance from the subject's ears, say). As a result, if topological spatial experience is possible, then it's possible to have a genuinely spatial experience without that experience being of something to the left or right, or in front of or behind, something else. To the extent that sensing movement requires sensing a change of locations, if in sensing one's conscious thoughts one senses topological properties only, then the only way that conscious thoughts could be experienced as moving would be for the thoughts to change their seeming location: from being seemingly inside (the subject's head) to being seemingly outside (out in the world), or vice versa. Such experience, of the thought seemingly leaving one's head seems to be what happens in thought broadcast. But such a spatial experience would not allow for the sense that one's thought could move around one's head. Its experienced spatial properties would not allow for such a fine-grained differentiation.

To this it might be objected that it's really hard to make sense of an experience representing topological properties only. Don't we always have a more detailed awareness of the spatial layout of the environment than just that some things are outside other things?

However, the possibility of topological perception in vision (the pickup of topological properties by a human visual system) is a meaningful empirical hypothesis. Some psychologists have tried to show that adults visually perceive topological relations (Chen,

1982). On the other hand, Jean Piaget and Barbel Inhelder (1963) famously proposed that as children we start out by perceiving topological relations, only later to develop the ability to perceive the environment in a somewhat less crude manner. These examples show that there's nothing absurd or incoherent about an experience of merely topological properties. I take it as uncontroversial that such an experience would be spatial, in virtue of representing spatial properties.

4.3 Introspective Evidence of Spatial Cognitive Phenomenology

The next sections are devoted to showing how the three lines of thought identified above collectively give dialectical advantage to the view that thoughts always have spatial phenomenology. I start with introspection.

Naïve introspection reveals that when I'm thinking about something consciously, the thoughts do seem to occur in my head. That is to say, there is a region of space, enclosed by the boundaries of my skull (roughly), where my conscious thinking seems to be coming from. That sense is especially vivid when I'm thinking in silent speech: producing a string of English sentences which seem to arise in my head, behind my eyes and between my ears. The sentences are not audible, but, with a little bit of focus, I could produce them with a different "accent" (or does that amount to imagining what I would sound like, were I raised in Yorkshire?) or in a different language.⁹ Beyond that, the thoughts (or, rather, their auditory vehicles) don't seem to have a very precise location, and naturally they don't seem to have any spatial extension. Neither does their location within the head seem to be capable

⁹ This phenomenological description does not commit me to any view about the connection between thought and language. It is just a description of how certain thoughts manifest themselves to my consciousness. But this is a distinct question from the one about the essence of thought in general. It might be that the strong connection between language and conscious thought is atypical of the connection between language and thought in general.

of changing (moving a bit towards the left ear, say). But I experience them to be in my head, for all that.

Obviously, mere introspection doesn't get us very far. For it is perfectly open to the opponents of spatial cognitive phenomenology to appeal to their own introspection and deny any phenomenal location properties in at least some conscious thoughts. Which is precisely what some philosophers have done (Bayne & Chalmers, 2003; Dainton, 2000; Roelofs, 2014; Watzl, 2014). This might be thought sufficient to support the view that some thoughts have, and others lack, spatial cognitive phenomenology. Introspective evidence, then, favors the moderate view mentioned above. This conclusion would have to rely on the assumption that the reports of both sides to the debate are equally good. But the assumption is not mandatory.

One worry to have here is a worry about the reliability of introspection as a source of evidence in general. If one were skeptical of introspection in general, the considerations of the previous section would be entirely irrelevant to the issue in question. If so, they would endow neither side with the dialectical advantage. If introspection is unreliable, then we have an impasse.

Another worry about the conclusion that introspective evidence favors the moderate view could be that introspection is simply powerless, on its own, to resolve this dispute (a situation analogous to the debate about cognitive phenomenology). While this approach does allow introspection to play some role in the debate, it also lets us treat the dialectical situation described in the preceding paragraphs as an impasse, rather than a clear victory for the moderates. Just as an introspective report of a conscious thought devoid of sensory phenomenology does not constitute a knockout argument against the view that conscious

thoughts always have sensory phenomenology, an introspective report of a conscious thought devoid of any spatial phenomenology doesn't constitute a knockdown argument against my view.

Lastly, it might be held that if the moderates were right that introspection supports their view, this would impose an additional burden on their theory, one that the strong cognitive spatiality thesis does not face. The burden is to explain why there's a phenomenological divide between the two kinds of conscious thoughts: those that have, and those that lack, spatial cognitive phenomenology. It seems intuitively more plausible to expect the phenomenology of thinking to be homogenous (for example, if conscious thinking correlates with the activation of the same neural structures, regardless of the content being thought, then it seems plausible that the same neural structures being activated would yield roughly similar phenomenology; alternatively, assuming that the phenomenological properties accrue to conscious thinking not only due to the content being thought, but also in virtue of the propositional attitude being consciously entertained, then it seems plausible that we will have broad phenomenological similarity among various conscious cognitive states). So it might be held that even though introspection favors the denial of the strong cognitive spatiality thesis, the moderate view incurs an explanatory debt elsewhere.

In none of the above cases, then, does it seem that the consideration of introspective data provides any of the views concerning spatial cognitive phenomenology with a clear advantage. This is why we need to base our arguments on more than introspective data (I don't think we can ever escape introspection entirely when debating such issues; which might be problematic if one is skeptical of introspection in general, but these skeptical worries would apply equally to proponents and deniers of spatial cognitive

phenomenology). The next two sections try to broaden the range of considerations in favor of thinking that the strong cognitive spatiality thesis enjoys a dialectical advantage, though ultimately they may fail to satisfy a hardline introspection skeptic.

Thus, in the next section I argue that the debate on cognitive phenomenology makes the spatial phenomenology view more appealing than its denial.

4.4 Cognitive Phenomenology and Spatial Cognitive Phenomenology

In this section, I propose that spatial phenomenology of thought sits more easily with the majority of the positions one can adopt with respect to the cognitive phenomenology debate in general. I set aside the view that there's nothing it is like to undergo thoughts and other cognitive states. The cognitive phenomenology debate, as I understand it, is concerned with whether conscious thought has phenomenal properties that are distinct from phenomenal properties associated with conscious sensory states and conscious bodily states. In what follows I adopt Charles Siewert's characterization of sensory phenomenal features as "those features whose possession is found in the activity of various standardly recognized perceptual modalities (vision, hearing, etc.) along with bodily feelings of pain and pleasure, cold and warmth, and kindred sensations, together with whatever analogs of these there might be in imagery" (2011).

Now, here are the options one can pick with respect to the phenomenal properties of conscious thoughts: (a) all cognitive phenomenology is sensory (all phenomenal cognitive states have only sensory phenomenal properties, including the phenomenology of proprioception: a view endorsed by e.g. Jesse Prinz (2011), Adam Pautz (2013), Peter Carruthers and Benedicte Veillete (2011) and Michael Tye and Briggs Wright (2011); (b) there is proprietary thought phenomenology but it is always associated with sensory

phenomenology (i.e. there are no wholly cognitive phenomenal states, but the cognitive states have some extra phenomenology that wholly sensory states lack; Elijah Chudnoff (2015) seems to favor this view); (c) some conscious thoughts may have proprietary cognitive phenomenology in the absence of any other phenomenology (i.e. there could be wholly cognitive phenomenal states; a view endorsed by, e.g., Russell Hurlburt (1997; 2008) and Uriah Kriegel (forthcoming)).

I take it that options (a) and (b) are probably the ones most easily accommodating of spatial cognitive phenomenology. Indeed, one may argue that views like these *entail* that cognitive phenomenology has a spatial aspect to it. One argument could go like this:

- (1) All conscious thought involves sensory phenomenology.
- (2) All sensory phenomenology has a spatial aspect to it.
- (3) So, all conscious thought phenomenology has a spatial aspect to it.

Premise (2) does require some defense. However the kinds of considerations in favor of its truth are not hard to find. At first sight, it does seem that when we sense things, we sense them to be out there or in here, at least. They are not sensed to be nowhere at all (it is hard to understand what this could even mean). It seems to me that less argumentation is required to show that (2) is true, than to show that the paradigmatic cases of conscious thoughts involving sensory phenomenology (thinking in words of one's language or entertaining mental images) lack a spatial aspect. It is difficult to imagine hearing silent speech without any sense whatsoever of where it was originating. Despite prominent views to the contrary (see e.g. (Strawson, 1959)), it appears phenomenologically apparent that auditory experiences have spatial phenomenology. James Van Cleve (2006), for instance, uses a simple appeal to phenomenology to dispute Strawson's claims, while Casey

O'Callaghan (2010) offers a critique of Strawson's arguments which, to my mind, is compelling. Van Cleve says: "Strawson [is] wrong to regard hearing as an aspatial sense. I would think [him] wrong on purely phenomenological grounds quite apart from physiology" and then explains in a footnote: "I hear the notes of the mockingbird outside my study as possessing an intrinsic over there-ness. I do not merely associate the melodious song of the mockingbird with what I shall see if I look to the left and the raucous call of the bluejay with what I shall see if I look to the right."

O'Callaghan raises similar worries for Strawson's view:

Location information detected by auditory processes is perceptually and consciously manifested. It is not merely delivered subconsciously to epistemic processes that manage belief formation. Though the mechanisms of auditory localization utilize temporal information, such as wave onset delays, and phase information, subconscious localization processes result in conscious perceptual access to locations. This is the experience of audible direction and distance from oneself. Support for this claim goes beyond introspection. Behavioral evidence from performance shows clearly that localization is a low-level, automatic perceptual process that does not require inference, effort, or extra-auditory cognition to "work out". Turning reactively to a sound behind the head is virtually unavoidable... Subjects who report and act upon the audible locations of events in egocentric space neither guess nor report beliefs lacking ground in auditory experience. It seems we hear where things are and where sounds come from ... [A]udible qualities seem to have audible locations—a fact demonstrated by behavioral and introspective tasks" (2010)

I admit to being convinced by this. Auditory phenomenology seems to be spatial phenomenology.

It is even more difficult to imagine "seeing" a mental image without it seeming to occupy some position in a visual field (intelligible questions could be raised about whether mental images are seen to occupy the same space as objects represented in perception; but they do seem to occupy some space or other).

So, with respect to options (a) and (b) the proponent of spatial phenomenology of thought is in a better position than the denier or the moderate. This is because there does not seem to be an obvious way of denying some spatial aspect to our sensory and proprioceptive phenomenally conscious states. If (a) or (b) are true, then the strong cognitive spatiality thesis, being entailed by them, inherits their modal force. Hence, if (a) or (b) are true, the moderate position is false (perhaps necessarily so).

What about option (c) then? Clearly, the worst case scenario for proponents of spatial cognitive phenomenology would be if the debate turned out in (c)'s favor. For then it would seem obvious that wholly cognitive phenomenal states need not involve spatial phenomenology. One could even try to argue that (c) entails that some conscious thoughts lack spatial phenomenology. To do it one would have to argue like this:

- (1) Only phenomenally conscious sensory states exhibit spatial phenomenology.
- (2) There could be wholly cognitive phenomenal states.
- (3) A wholly cognitive phenomenal state is not a sensory state.
- (4) So, there could be phenomenal states that lack spatial phenomenology.

But, fortunately, (1) doesn't appear to be true: one can feel pride in one's chest (in the spatial sense of "in"), desire in one's loins and anxiety in one's stomach. It's at least not obvious that these states have a sensory component, and they seem to have a spatial component.

One can still say that the states mentioned are body-involving, and having to do with bodily processes, and that qualifies them as sensory, in contrast to cognitive phenomenology. But if one is so permissive with assigning body-involving phenomenology to various (not *obviously* sensory or proprioceptive) conscious states, then

it seems to beg the question against the proponent of spatial cognitive phenomenology to insist that cognitive phenomenology lacks any spatial aspects. We'd need a principled reason to exclude cognitive phenomenology from the class of conscious states that could have spatial aspects. But the argument was intended to provide such a principled reason! So it seems circular.

It is, moreover, not obvious why it would be part of the essence of wholly cognitive states that they couldn't be felt by the subject to occur in some part of his body, or inside him in general. If it isn't part of their essence, then spatial cognitive phenomenology is compatible with (c). After all, we do have ways of representing spatial properties in a non-sensory way (think e.g. about representing a sphere by the equation " $x^2+y^2+z^2-1=0$," or representing multidimensional spaces that seem impossible to visualize and perceive).

In sum, at least this asymmetry seems to present itself when it comes to the issue of spatial cognitive phenomenology in the context of the debate on cognitive phenomenology in general. The positions that are most friendly to spatial cognitive phenomenology seem to entail it, given some plausible assumptions. The position least friendly to it still seems compatible with spatial cognitive phenomenology. On the other hand, it appears that the denial of spatial cognitive phenomenology is incompatible with (a) and (b). And (c) does not appear to favor the denial of spatial cognitive phenomenology over its opponent (at least not to the same degree that (a) and (b) favor spatial cognitive phenomenology over the moderate view).

All in all, then, it seems now that spatial phenomenology of thoughts has a better chance of being accommodated by a wider variety of positions within the cognitive phenomenology debate than the denial of spatial phenomenology to thoughts. This makes

the former view *prima facie* preferable. Given that introspective data did not give any position a clear advantage, the strong cognitive spatiality view seems to be preferable on the balance of all evidence so far considered.

4.5 Thought Broadcast and Spatial Cognitive Phenomenology

I will now turn to the third way of supporting the strong cognitive spatiality view, based on thought broadcast – a pathology of cognitive phenomenology that occurs in some schizophrenic patients. I will argue that the most plausible explanation of the schizophrenic patients' phenomenology favors the postulation of spatial cognitive phenomenology over its denial.

4.5.1 Thought Broadcast

In general, patients suffering from thought broadcast feel that their thoughts escape their minds, and may then become available for others to access. I will argue that the explanation of the phenomenology of this symptom of schizophrenia favors the postulation of spatial cognitive phenomenology over its denial.

Several related but distinct phenomena have been described as thought broadcast in the literature. The name is somewhat misleading (“thought diffusion,” which was originally used to describe this phenomenon would perhaps be more apt, but that name is no longer current in the literature, so I will stick to the standard label) in that not all clinical definitions of thought broadcast require that the thought's content actually become available (i.e. broadcasted) to others. It is these latter cases (thought broadcast, so to speak, without its content becoming available to others) that are of particular interest for my purposes. This type of thought broadcast is described in the literature in the following ways:

Diffusion or broadcasting of thoughts. The patient, during the process of thinking, has the experience that his thoughts are not contained within his own mind. The thoughts escape from the confines of the self into the external world, where they may be experienced by all around. There is usually a secondary delusional explanation of this phenomenon which may invoke the use of telepathy, television, etc. (Mellor, 1970)

Thought broadcasting. The patient's experience that as his thoughts occur they are escaping from his head aloud into the external world. (Taylor & Heiser, 1971)

[T]he subject is quite certain of “negatively” being aware that he has lost HIS OWN thoughts, feelings and so on because in some way they passively diffuse into or are lost to the outside world against his will (e.g. thought broadcast). (Koehler, 1979)

[T]hought broadcast ... (ii) wide: either thoughts leave the confines of the subject's mind but are not shared or they are so loud that others can hear them. (O'Grady, 1990)

[In thought broadcast t]he patient experiences his thoughts as escaping silently; they may or may not be available to other people. This definition does not stipulate that thoughts become audible. They are perceived as escaping silently, with no mechanism being specified as to how they escape. (Pawar & Spence, 2003)

Mellor also offers an actual report of a patient with thought broadcast: “A 21-year-old student said: ‘As I think, my thoughts leave my head on a type of mental ticker-tape. Everyone around has only to pass the tape through their mind and they know my thoughts.’” (1970).

Thought broadcast, then, involves (at times, at least) the sense that one's thoughts leave one's head. To account for such cases, we seem to need to appeal to the spatial phenomenology of thoughts. The next section explains why.

4.5.2 Spatial Phenomenology in Thought Broadcast

What goes wrong in these cases of thought broadcast? What disruption of ordinary phenomenology do the above descriptions convey?

In my view, the best explanation of thought broadcast phenomena involves an appeal to spatial phenomenology: what goes wrong in thought broadcast is that the patients feel

their thoughts to change their felt spatial location. Most other aspects of having conscious thoughts seem to remain the same in broadcasted and regular thoughts (cf. e.g. Bortolotti and Broome (2009), Martin and Pacherie (2013)). Consider: the patients seem to have good, immediate, first-personal access to their thoughts (so immediacy is preserved); the thoughts need not be accessed by others (so privacy is preserved); and they seem to be both owned and authored by their subjects (so both the sense of ownership and the sense of agency are preserved). Suppose this is right. Then, the only thing that's odd about these thoughts is that they seem to occur at some unusual location, i.e. outside the subject's head!

So, rather than occurring in the head, as we normally feel our thoughts to occur, the thoughts are felt to escape, or leave the subjects' heads and become lodged, or perhaps spread out, outside their bodies. The thought is sensed to have changed its topological properties with respect to the subject (or the subject's head). The thought has ceased to occupy one location, and began occupying another location. It is quite natural, I submit, to explain thought broadcast in this way. Yet, this natural interpretation of the phenomena is not available to the (wholesale) denier of spatial cognitive phenomenology. And it's only available to the proponent of the moderate view if she is willing to accept some extra *ad hoc* hypotheses as well (I will explain more about this later on).

It might be argued in reply that the pathology involved in thought broadcast consists in the spatialization of conscious thoughts rather than a disruption of the normal spatial 'feel' conscious thoughts have. That is to say, what accounts for thought broadcast is not that the patients begin to feel their thoughts to be located outside their heads, rather than inside it as they normally do, but rather that they begin to feel them to be located at all.

One reason to resist this line of thought is this: if the above is right, then it's quite surprising why the patients go to the trouble of saying that thoughts *leave* or escape their heads, rather than settle for something like: my thoughts are out there, floating about etc. The contrast alluded to does not seem to be between lack of spatial properties and presence of spatial properties. Rather, it is a contrast between a thought being inside and its being outside; it is the unusual felt location that accounts for some cases of thought broadcast, not the fact that the thoughts seem to have a location at all (if it were, one would expect the patients to remark on this new aspect of their phenomenology; as far as I can tell, they don't). Similarly, there are pathological cases in which patients ascribe spatial properties to conscious thoughts that do not seem to be part of ordinary cognitive phenomenology. Some patients seem to feel their thoughts to be to the left or right of one another, e.g., or to press against a side of their skull—cf. (Parnas et al., 2005). It appears that in such cases it is likewise the disruption of the normal spatial feel of the thought that is responsible for the reports, and not the sudden acquisition of spatial phenomenology. The reason is the same as with thought broadcast. The patients don't remark on the very existence of felt spatiality of the thoughts as odd. Rather, they remark on the *disrupted* spatiality as odd.

Additionally, if all cases of thought broadcast involved was the loss of privacy to one's thought, it would be surprising why the patients would persist in using spatial vocabulary when describing their symptoms. However, if those cases (referred to in the literature as “narrow” thought broadcast) were the only cases of thought broadcast, the denier of spatial cognitive phenomenology could say this: what the patients experience is that their thoughts' contents are available to others. Undergoing such an unusual experience requires an explanation. The patients explain their experience of their own thoughts' being available

to others by saying that their thoughts leave their own heads. Yet, this need not be seen as a *description* of their phenomenology, but rather as a quasi-theoretical *explanation* of the patient's seeming to lose exclusive access to their thoughts.

The above line of thought strikes me as plausible. So it is fortunate that there is some reason to believe that "wide" symptoms of thought broadcast really occur. O'Grady and Pawar and Spence both present data according to which some schizophrenics (and a small number of non-schizophrenics) are reported to exhibit "wide" thought broadcast.

There are two further worries one could raise at this point. First, the moderate could accept that some thoughts (in particular, the thoughts of the thought broadcast subjects) involve spatial phenomenology but deny that the results generalize to *all* instances of conscious thinking. Second, one could say that the patients just happen to be ones whose thoughts have spatial phenomenology, but there's no reason to generalize to all *subjects* of conscious thinking. These seem to be *ad hoc* responses, however. What, exactly, is the reason that we should not generalize to all subjects or to all thoughts?

It might nevertheless be granted that generalizing to all conscious thought subjects from the few cases of thought broadcast is too hasty. Even so, the burden now is firmly shifted onto the moderates and the deniers of spatial phenomenology to provide reasons why thought broadcast cases are not illustrative of disruptions to certain aspects of everyday cognitive phenomenology. In the absence of such reasons, it seems to me that we can adopt the spatial cognitive phenomenology view as default.

In sum, considerations concerning thought broadcast suggest that spatial cognitive phenomenology is a useful explanatory resource, and that the deniers (even those of the moderate variety) may face some dialectical disadvantages when trying to explain these

phenomena. Additionally, absent a principled reason to believe this, it seems completely *ad hoc* to say that spatial cognitive phenomenology occurs in the pathological cases but does not generalize to all cases of cognitive phenomenology.

4.5.3 A Possible Objection

To all this the following objection might be raised: Some cultures seem to think that the heart is where the soul resides and thinking originates. It is parochial to think that all thought phenomenology involves the sense of thought occurring in one's head. It's not likely that such phenomenology would be avowed by, say, Aristotle or Mencius.

Suppose it is replied that people feel the thoughts merely to be located inside their boundaries, and cultural variation plays a role only when it comes to specifying the more exact location of the thought (head, heart etc.). But if the friend of spatial phenomenology could allow for cultural variation (presumably in the form of culturally shared beliefs and naïve theories) to influence some spatial phenomenology of thought (or descriptions of such phenomenology), then what's stopping a denier from going all the way and attributing *all* spatial talk to culturally influenced quasi-theories about where thinking occurs? This objection may be formulated as follows:

It might be that all this "in the head" talk is just a reflection of a naïve theory about where the thinking machinery of human beings is located. Many of us believe that the brain is the organ of thought and that this organ is in the head. We shouldn't read more into the speech patterns of schizophrenic patients and ordinary subjects than this. In particular, we shouldn't infer that there is any spatial phenomenology captured by spatial idioms. Rather, there is just cognitive phenomenology whose description people embellish with a quasi-theoretical addition. Hence, people who say things like "there's a thought in my head"

make an inferential step from “I’m conscious of having a thought” via “thinking is caused by the brain” and “my brain is in my head” to “there’s a thought in my head.” But only “I’m conscious of having a thought” should be read as a phenomenological report. The rest is theoretical superstructure.

I have the following to say in reply: At least some philosophers have endorsed the view that their conscious thoughts seem to exhibit spatial phenomenology (e.g. Daniel Dennett in “Where Am I?” (1981) says “when I thought “Here I am,” *where* the thought occurred to me was *here*, outside the vat, where I, Dennett, was standing staring at my brain [first italics mine]”). Should we assume that they too are bewitched by the metaphor and the naïve theory of where the organ of thought is located in the human body? It is, of course, not impossible that they are. But a theory that explains away introspective judgments by accusing those avowing them of confusion has to rely on more than mere introspective judgment to argue its own case. For one could equally well argue that the non-spatial view is an effect of being bewitched by language (or a quasi-theory) that treats minds as fundamentally different from bodies. Our different takes on naïve metaphysics may pull our descriptions of cognitive phenomenology in different directions. So what? As things stand, the deniers of spatial cognitive phenomenology use appeal to introspection as their main strategy. And this, it seems to me, provides them with no dialectical advantage if all they can do to discredit the opponents’ view is to explain away *their* introspective descriptions. This is because such explaining away is a double-edged sword, and it cuts both parties to the debate in equal measure.

A subtler explanation of why there is a divergence in introspective reports is in terms of cognitive penetration. It might be said that the belief about where the thought originates

cognitively penetrates thought phenomenology. That is, the phenomenology of thought is causally affected by the belief concerning the thought's origin in such a way as to yield spatial phenomenology: the genuinely felt spatial location of the thought. However, where the belief is absent, the phenomenology lacks this spatial element. This explains intersubjective variation.

The extent of cognitive penetration, and its very existence, are subjects of much controversy. But even assuming that cognitive penetration is real, it would be too hasty to conclude that it provides an explanation in this sort of case. This is because an appeal to cognitive penetration would not be in a position to explain the variation of introspective reports on cognitive phenomenology within a single subject (and Dainton and Roelofs seem open to admitting that while some (or even most) thoughts have spatial phenomenology, others lack it; that is, they seem open to the view that both kinds of thoughts seem able to occur within a single subject). It would be extremely puzzling to maintain that my belief that my thoughts originate in my head should penetrate some, but not all, of my cognitive phenomenology.

As a result, the appeal to the background theory of the location of our thinking machinery does not appear to succeed in blunting the force of the considerations advanced here. Regardless of whether the objection to the strong cognitive spatiality thesis has it that the background theory informs introspective judgments without affecting phenomenology, or whether the background theory is supposed to cognitively penetrate the phenomenology itself, the objection doesn't succeed. This is because either we can explain away introspective judgments of the proponents of the weaker theses concerning the thoughts' felt spatiality (just as they claim to explain away judgments supporting the strong thesis),

or we may doubt the validity of appeal to cognitive penetration as an explanatory tool in the case of thought phenomenology. However, if we can explain away introspective judgments of the proponents of the weaker thesis concerning thought spatiality, the other side is left with nothing else to appeal to. In contrast, the fans of the stronger thesis may appeal to considerations from beyond introspection to argue their case. This gives the stronger thesis the dialectical advantage.

4.6 Summary

I have surveyed three different areas of theorizing about thought phenomenology (naïve introspection, the cognitive phenomenology debate, and some deviant cases of thought phenomenology) to conclude that it is *prima facie* more reasonable to adopt spatial cognitive phenomenology as a default position. This shifts the dialectical burden onto the denier of spatial phenomenology. I have also examined some objections to the spatial cognitive phenomenology claim, and tried to show how they fail. Given all that, I think it is reasonable at this point to assume that our thoughts involve spatial phenomenology.

The spatial properties involved in thought phenomenology are of such nature that cognitive experiences are spatially unified with other spatial experiences of the subject. This is because cognitive phenomenology seems always to involve, at least, localizing the thought (thinking activity) either inside or outside the subject, i.e. either within or outside the subject's boundaries. But these boundaries are experienced as spatially related to other things sensed and perceived. Hence, the location of a thought is experienced as belonging to the same space as the felt locations outside and inside the subject.

Chapter 5: What makes up a Mood Experience?

5.1 Moods and Their Characteristics

In this chapter I argue that the phenomenal character of a mood experience (that is, what it's like to be in a mood) wholly depends on affective modifications (appropriate for the mood in question) to the phenomenal characters of one's non-mood experiences. I argue that this view accounts for all distinctive aspects of mood phenomenology. I also point out what consequences this view has for using mood experiences as counterexamples to the Intentionalist project of reducing the phenomenal character of experience to its intentional content. Lastly, I explain how this account is consistent with the thesis that spatial unity is necessary for phenomenal unity.

I will start by explaining what I mean by mood experiences and what I take to be the characteristics of such experiences that every theory of mood experience ought to account for. I will also briefly introduce existing theories alongside my own theory of mood experience. Then I will explicate my own theory and show how it explains every characteristic feature of mood phenomenology. I will end with some remarks on the relevance of my view for the debate on Intentionalism in general and the account of the unity of consciousness I'm offering in this work.

Paradigmatic cases of moods include elation, depression, and anxiety. Mood experiences are generally agreed to have three characteristic features: (a) mood experiences are phenomenologically similar to corresponding emotional experiences (e.g. anxiety is similar to fear, and elation to happiness); (b) mood experiences, in contrast to emotions, typically seem to have unspecified objects, or even seem to lack intentional objects

altogether (they don't seem to be about anything); and lastly, (c) they also have far-reaching effects on other experiences, and indeed, on our entire psychological life.

I take it that while (a) and (b) are intuitively understandable, (c) requires additional elucidation. John Haugeland expresses this last feature of moods eloquently when he writes:

moods are pervasive and all-encompassing ... The change from being cheerful to being melancholy is much more thorough and far-reaching than that from having a painless foot to having a foot that hurts. Not only does your foot seem different, but *everything you encounter seems different. The whole world and everything in it, past, present, and future, becomes grayer, duller, less livable. Minor irritations and failings are more conspicuous and less remediable; ordinary things are no longer fun, lovely, or pleasing ... moods not only affect how things look, they affect how one thinks.* What seems reasonable when you're cheerful seems foolish when you're melancholy, and vice versa. Likelihoods and improbabilities invert, as do what seems relevant to an issue and what seems beside the point (Haugeland, 1978; italics mine).

Since (a)—(c) are commonly agreed to be characteristic features of the phenomenal character of moods, in what follows, I will use the term 'mood experiences' to refer to phenomenally conscious experiences that have these three phenomenal features.

Below I will argue for the view that the phenomenal character of a mood experience wholly depends on what it's like to undergo other, non-mood experiences modified in a certain patterned way and that this explains why moods have the features they do.

5.1.1 Dimensions of Mood Theories

In this subsection, I introduce two dimensions along which theories of mood experiences may vary. First, theories of mood experience typically take a stand on whether mood experiences have *intentional content*. Thus, we have intentionalist accounts of mood phenomenology (Kenny, 1963; Mendelovici, 2014; Price, 2006; Solomon, 1993; Tye, 1995, 2008) according to which mood experiences have such content, and non-

intentionalist ones (Russell, 2003; Searle, 1983) according to which mood experiences are not about anything.

Second, theories of moods vary according to whether they take mood experiences to have the phenomenology that's independent from the phenomenal character of other experiences, or whether they take mood phenomenology to be so dependent. Let's clarify these notions a little.

One example of a dependent experience is an experience which attributes properties detectable by different senses (sound and color, e.g.) to the same object. Take an audiovisual experience as of a moving car. That experience's phenomenal character (of a red oblong car making a loud noise, say) wholly depends on the phenomenal character of the visual experience (seeing the car's color and shape) and the phenomenal character of the auditory experience (hearing the sounds the car is making). A dependent experience is one whose phenomenal character wholly depends on the phenomenal characters of other experiences. That is, it couldn't occur in the absence of the experiences it is dependent on. Let us call any experience that is not dependent, a *basic* experience.

I call the views according to which mood experiences are basic, the Ontological Independence of Moods views (OIMs), and the views that deny it, the Ontological Dependence of Moods views (ODMs).

In the next section I divide the accounts of mood experiences by how they address the issues of: (1) intentionality and (2) ontological independence of mood experiences.

5.1.2 Classifying Mood Theories

There are multiple accounts of the phenomenal character of mood experiences. It is variously taken to consist in: raw feels (e.g. James Russell); representations of unattached

affective properties (e.g. Angela Mendelovici); representations of one's overall condition (e.g. Michael Tye); global representations of the world as a whole (e.g. Robert Solomon, Carolyn Price) or of a series of objects (e.g. Anthony Kenny). I will briefly introduce these views via the way each addresses the two dimensions mentioned above.

On *Russell's* view, for a subject to undergo a mood experience is for the subject to experience a raw feel: a purely qualitative, free-floating conscious sensation that is not directed at any object. E.g., when a subject experiences anxiety, she has a mental state essentially consisting in some sort of non-representational anxiety quale. The subject just feels anxious.

The phenomenal character of the raw feel experience does not appear to depend on other non-mood experiences. It seems possible to experience a raw feel independently of any other experiences you're having. Hence, classifying Russell's view as OIM appears appropriate.

On *Mendelovici's* view a subject is in a mood when she *represents* an unattached affective property. E.g. when a subject experiences anxiety, she has a mental state essentially consisting in representing the unattached property of scariness.

This view is also, I take it, an OIM. This is because the phenomenal character of the representation of an unattached property does not seem to depend on other non-mood experiences.

Tye's one view is that a mood experience essentially consists in having an experience representing the overall state of oneself, in something like this way: "[i]f one feels elated, one experiences a change in oneself overall. The qualities of which one is directly aware in attending to how one feels on such an occasion are experienced as qualities of oneself.

One is aware of a general sense of buoyancy, of quickened reactions, of somehow being more alive” (2000). If one reads these remarks as saying what a mood experience is, then Tye’s view counts as OIM. Being directly aware of my own qualities seems independent of any other experiences.

Another view of Tye’s (advanced in his (1995)) is that mood experiences represent departures from the regular range of one’s body’s physical states. In other words, moods represent our bodies in a certain way. This representation, again, doesn’t seem to depend on any other experience (it might depend on a generalized representation of my body, but that representation need not be itself conscious). Both Tye’s views, then, are OIM.

The main motivation for the non-intentionalist OIM is phenomenological: mood experiences don’t seem, from the inside, to be about anything in particular. When you’re elated or depressed, there need be nothing about which you feel in this way. That is, mood experiences introspectively don’t seem to be intentional states. The main motivation for intentionalist OIM (as it is for intentionalist views about moods in general) is theoretical. The idea is to save Intentionalism (the theory that phenomenal character reduces to intentional content) from the apparent counterexample in the form of a mood experience that is phenomenally conscious without being intentional.

The accounts according to which moods represent *the world as a whole* can be given both an OIM reading and an ODM reading. That is, the mood experience can be a basic experience (basic in virtue of having its own independent phenomenology) representing the whole world, or the phenomenal character of the representation of the whole world may be dependent on what it’s like to consciously represent parts of the world. On the first reading, the basic experience could be a single thought with the content “The world as a

whole is wonderful” (see e.g. (Tye, 2008)). You may consciously entertain this thought in the absence of any other experience. So its phenomenal character doesn’t depend on other experiences.

Alternatively you might think that the representation of the world as a whole is, typically, perceptual and hence, typically, multimodal (a version of this view is developed by Price). On such a view the phenomenal character of the mood experience would depend on representing parts of the world in a certain way. At any rate, the world-as-a-whole intentionalist views about mood experiences could be either OIM or ODM.

Lastly, consider *Kenny’s* view according to which moods represent a series of objects. E.g., the objects of a depressed mood “are the things that seem black.” It also seems that this is an ODM view. Your mood experience’s phenomenal character depends on experiencing particular objects. These multiple objects, according to Kenny, are the intentional objects of a mood experience. The generalized conscious representation that constitutes the mood experience (“objects A—N seem black”) depends on the conscious representations of particular objects (“A seems black, ..., and N seems black”).

And, finally, this is the view I will argue for:

The Patterned Modification View (henceforth, PMV): What it’s like for the subject to undergo a mood experience wholly depends on what it’s like for the subject to have non-mood experiences affectively modified in a certain patterned way appropriate for the mood in question.

My view is obviously an ODM view. A mood experience cannot occur without those other experiences obtaining. However, in contrast to the existing accounts I have presented, PMV leaves the question of mood intentionality open. It will, I believe, emerge in the

course of this discussion that I have a good reason for abstaining here. Put simply, it's because other difficult questions need to be answered before we can say whether mood experiences are intentional states or not.

To articulate and motivate PMV, in Section 5.2 I will say more about modification of mental states in general, and the modifications relevant for mood experiences in particular; Section 5.3 explains the notion of pattern in more detail, in preparation for the arguments of Section 5.4 where I will argue that a mood experience cannot occur in the absence of the pattern of modification, and Section 5.5 where I argue that PMV explains all the characteristic features of mood phenomenology; Section 5.6 briefly notes the relevance of these considerations for that part of the debate on Intentionalism that appeals to moods and Section 5.7 relates the considerations of this chapter to the whole project.

5.2 Modification

In this section I introduce and explicate the notions of modification, and the close connection between mood experiences and modifications of non-mood experiences. I will proceed by listing examples of modification of mental states outside the context of a mood, and then I will point out some general characteristics of such modifications. After that, I will list some examples of modifications by moods.

5.2.1 Examples of Modification in General

Suppose you're reading a detective novel. About halfway through, you begin to suspect who the kidnapper is. You believe, with the credence of .6, say, that it's the bodyguard. As you read further and more clues are uncovered, you acquire more and more certainty that the bodyguard is, indeed, the kidnapper. Your credence rises to, say, .9. Your belief, then, has been modified along the credence *parameter*. It has been modified by the evidence you

acquired. Furthermore, if the belief were conscious, then there would be a phenomenal difference between consciously believing with .6 credence that the bodyguard is the kidnapper, and consciously believing with .9 credence that the bodyguard is the kidnapper. This is our first example of modification.

Perceptual states may also be candidates for modification. One oft-cited example is that hills look steeper when you're wearing a heavy backpack. Suppose this is right (though Chaz Firestone (2013) offers a powerful critique of these sorts of cases). Looking at some hill without wearing a backpack, you experience it as steep to degree N . Once you put on the backpack, however, when looking at the same hill, in the same viewing conditions, you see it as steep to degree $M > N$. In such a case we might say that your perceptual experience of a hill has been modified along the parameter of perceived slope.

These are, we might say, single, linear parameter modifications. But often isolating just one parameter that gets changed will be difficult to do. Consider phenomenal contrasts between an expert and a novice (often appealed to in the debate on high-level properties, most famously by Susanna Siegel (2011)), as outlined by Kevin Connolly (2014): "Due to practice or experience with properties in their expert domain, experts in some fields perceive the world differently from non-experts. Cabernet Sauvignon tastes different to a wine connoisseur than to a novice. Beethoven's Ninth Symphony sounds different to a conductor than to an untrained listener. What it is like to see a wren is different for an expert birdwatcher than a non-expert." Tasting Cabernet Sauvignon, at some time t , I experience its flavor in just the way a novice would. Then I undergo months of rigorous sommelier training. After the training period, at some time $t+1$, when I taste the same wine, in identical tasting conditions, I experience its flavor differently. In such a case we might

say that my experience of the wine has been modified along some parameter having to do with the flavor (complexity? Variety?). But it's not easy to specify informatively what aspect of the experience it is that has been modified. Intuitively, however, these examples belong together with the examples of modification in the belief case and the hill case. So I will not treat the difficulty in specifying the parameter along which the phenomenal change occurs as indicating some deep divide between cases of modification of an experience.

Lastly, consider the following case. Suppose that over the course of several months, I develop a distinctive liking for Sumatran coffee. Before, its taste was neutral to me. I didn't get any pleasant sensations when drinking it. It felt no different than drinking any other coffee, as far as the pleasure it brought me was concerned. But after a while, drinking Sumatran coffee started bringing with it not just the usual caffeine rush, but also a distinctive pleasure. The coffee began tasting more pleasantly than other varieties. Here, the parameter that accounts for the change is the degree of the pleasantness of experience.

I take the above cases as paradigms of modification. The perceptual ones seem to have at least the following in common: take a subject's experience E_1 . While keeping the distal stimulus, the external conditions, the subject's sensory apparatus (in general, as much as possible) constant, vary the experience along some parameter that yields a change in E_1 , so that there's a phenomenal difference between the new experience E_1' and the old experience E_1 , and had the subject not been in E_1 at t , he would not be in E_1' at $t+1$, and E_1 and E_1' score highly on any relevant measure of similarity. On some views, under these conditions E_1' would be a modified version of E_1 . On other views (e.g. if you thought that the identity of an experience is wholly determined by its phenomenal character), they

would be two distinct experiences. But whatever relation holds between E1 and E1', I take that relation to be crucial to my notion of modification.

Likewise in the conscious belief case, the intentional object of the experience, the mode of accessing that object (i.e. the propositional attitude), even the content of the experience ('that x is the kidnapper') remain unchanged. And yet, the difference in credences could make a phenomenal difference. So, the later conscious thought seems to stand in roughly the same relation to the earlier thought as E1' stands to E1.

5.2.2 Modifications and Moods

Moods are often noted, both in psychological and philosophical literature, for the effects they exert on the subject's entire mental life. But the way these effects are specified bears a striking similarity to how phenomenal contrasts exemplified above are introduced. Below is a sample of quotations displaying this tendency (italics mine).

[B]eing in a mood seems to involve profound and far-reaching effects on one's mental condition. The change from a depressed to a euphoric mood is much more profound than the change (say) from hoping that Tom will get the job to fearing that Tom will get the job. ... The most common metaphors [to describe this feature of moods] are from optics. One reads that moods are 'filters' which '*color perceptions*'. One is told of the 'rose-colored glasses' of euphoria and the 'rosy light of happiness'. Furthermore, one is informed that in euphoria and depression '*things look rosy or black*'. And from Solomon one gathers that certain moods '[cast] happy glows or somber shadows' on every object of experience. (Lormand, 1985)

moods are pervasive and global, having the capability of influencing a broad range of thought processes and behavior ... they can ... subtly insinuate themselves into our lives, *influencing what we remember of the past, perceive in the present, and expect from the future* (Morris & Schnurr, 1989).

When people are in good moods, they are thought to be 'wearing rose-colored glasses,' tending to *notice and interpret events in an overly positive light* ... When in bad moods, however, *everything seems hopeless* and people tend to *notice and interpret events negatively* (Rusting, 1998).

When I enter into a particular mood, for example when I become irritated, *I begin to see things differently — I begin to see objects in the world as annoying and malign* (Fish, 2005).

As [a subject] walks along outside, she encounters an ever-changing landscape of sights, sounds, and smells. Today, given her mood, the storefront across the way doesn't *seem nearly as shabby as it usually does*; the face of her taciturn neighbor *doesn't strike her as quite as menacing*; the car horns fail to disturb her and likewise for the odors emanating from the overflowing dumpster at the curb. Today, *everything she passes seems wonderful and new* (Kind, 2014).

Taking at face value the phenomenological reports articulated above, it appears that mood experiences do indeed involve the kinds of modification that the previous subsection talked about. When you're in a mood, the phenomenal character of your non-mood experiences (perceptions, bodily sensations etc.) is changed in a specific way. There is a clear phenomenal difference between how things strike us when we're elated (as wonderful, say) and how they strike us when we're depressed (as bleak and hopeless, say). Our experiences get modified, perhaps along such parameters as valence (seeing things in a positive vs. negative light) and arousal, and maybe some others, in such a way as to yield a phenomenal difference even though the moody and the non-moody experiences are as similar as possible in all other respects.

The above quotations can, as a first pass, be all said to express, in their various ways, the following thesis:

MANIFESTATION: Mood experiences manifest themselves in consciousness via modifications to the phenomenal characters of non-mood experiences.

I assume that MANIFESTATION is true. Moreover, since its truth is widely acknowledged, it should be accounted for by any theory of mood experience.

5.2.3 Beyond Modifications

More needs to be said about mood experiences than what MANIFESTATION says. Most importantly, it's not just one or two phenomenal states that need to be modified for you to have a mood-experience; rather, a significant number of phenomenal states needs to be modified to make it reasonable to speak of a person as experiencing a mood. This is made clear by the observation that moods have a 'global,' or 'pervasive' character. They affect all (or, more plausibly, a significant chunk of) your mental states (look back again at some of the descriptions cited above: when in a mood, *a lot* seems different to you). It would be very strange to say that you're in a mood when only a small proportion of your experiences is affectively modified (you're not really anxious when only one thing in the world seems scary to you). We might say that the mood experiences' global character is realized at the phenomenal level via the mechanism specified by MANIFESTATION.

So, when you're having a mood experience, many of your mental states are modified in a certain way. It is, however, implausible to say that it has to be *all* of your mental states that are somehow individually modified by the mood for you to have a mood experience. For suppose you're depressed. Everything in the world seems awful and pointless. And then you feel an itch (or a random thought occurs to you). Such an itch (or thought) could be affectively neutral. But it's implausible to say that you have thereby snapped out of your depression as soon as you felt it. This seems to suggest that while mood experiences are indeed global—that is, they involve many of the subject's mental states—they don't have to be *total*; that is, they don't have to involve all of the subject's mental states.

Lastly, we must note that not just any global modifications of experiences' phenomenal characters will do to count as a mood experience. There has to be some coherence to *how* your experiences are modified. For instance, if your neighbor seems more terrifying to you,

the dinner tastes more wonderfully than usual and the TV show appears boring and dull, you're not in any recognizable mood. Your affect is all over the place. On the other hand, if all those things seem dull, or wonderful, or menacing, then you are in a mood. So, not only do most of your mental states need to be modified for you to count as having a mood experience. In addition, the modifications need to exhibit a certain *orderly arrangement*. As Michael Tye (2008) remarks, describing what it's like to feel elated, "as my experience of elation continues, its object changes. But even though there is no single object, each object is represented *in the same way* – as wonderful." We may therefore conclude that modifications of your non-mood experiences need to conform to some *pattern* for you to undergo a mood experience. Below I present one way in which we could think more precisely about such patterns. It should be noted, however, that I intend what follows merely as a heuristic tool geared towards helping to understand moods in the way I see them. The slightly more formal treatment of the notion of a pattern that I present merely puts some of the characteristics of patterns in sharper relief than an intuitive understanding appealed to so far. However, I take it that these characteristics are already part of our intuitive concept of the pattern. Thinking mathematically about them simply brings out these characteristics more clearly.

5.3 Patterns of Modification

To understand what a pattern is (mathematically speaking) we have to start with the mathematical definition of randomness: "A series of numbers is random if the smallest algorithm capable of specifying it to a computer has about the same number of bits of information as the series itself" (Chaitin, 1975). If a series is not random, then the information about it can be compressed. In the remainder of this paper, I follow Daniel

Dennett (1991) and interpret the definition as a bi-conditional. Hence, if a series can be compressed, it is not random, i.e. it's patterned. For example, consider the series 010101010101. You can give the computer the instruction to print the series by instructing it to print '01' six times. That information is compressed; it contains fewer bits than the original. On the other hand, consider the series 18656773652090. There is (I think) no shorter way of giving printing instructions to the computer other than specifying the whole sequence verbatim.

Loosely applied to mood experiences, the idea is this: when you want to describe to a computer (or someone else) how different phenomenal states of a person (call him Bert) having a mood experience are modified, you can simply identify each of Bert's conscious states, and then identify the way in which that state's phenomenal properties are affectively modified. That is, you can tell the computer that Bert's experience is such that "The whole world and everything in it, past, present, and future, becomes grayer, duller, less livable. Minor irritations and failings are more conspicuous and less remediable; ordinary things are no longer fun, lovely, or pleasing." This is just a non-compressed list of the modifications to various experiences of the world that Bert is undergoing. But you don't have to specify all the modified phenomenal properties of Bert's every mental state to let someone else know what Bert's experience feels like. You can simply tell them that Bert is experiencing a melancholy mood.¹⁰ Thus, when a subject is undergoing a mood experience, the way his experiences' phenomenal characters change can be specified without listing all the experiences and all the modifications. This is why, in saying that

¹⁰ This, obviously, is an idealization, but it seems to point in the right direction.

someone is having a mood experience, we capture a certain *pattern* of affective experiential modifications.

So, when you're having a mood experience, a significant portion of your conscious states is modified in that your conscious experiences change in a patterned way. Now's the time to show that a mood experience depends on such a patterned modification and that this fully explains mood phenomenology.

5.4 Mood Experiences Depend on Patterned Modifications

I will argue here that no mood experience can occur without a patterned modification of non-mood experiences. Before that, I will show that the OIM positions are mistaken if they take the basic experiences to be sufficient for a mood experience. It will then turn out that considerations exactly similar to those that show the insufficiency of a basic experience can be harnessed to show the moods' dependence on patterned modifications.

Case 1: Basic experiences are not sufficient for a mood experience.

Suppose that all you need for a mood experience is some basic experience (a raw feel, or a representation of an unattached property, or a thought that everything is awesome). That is, once you have the appropriate basic experience, you're undergoing the corresponding mood-experience. If this is right, then, whatever your other mental states feel like, as long as you have, say, a melancholia raw feel, you're experiencing melancholia. However, this entails that it's possible for you to have a melancholy mood experience while all, or the majority of, your mental states are modified in a way typical of elation. So, it could be that you are in a melancholy mood while everything seems new and wonderful to you (and, moreover, *nothing* seems sad or cheerless!). This is counterintuitive and *prima facie* inconsistent with

the moods experiences' global character. But it seems to follow from the view that all you need to have a mood experience is to have the appropriate basic experience. Hence, I conclude that just having a basic experience of some sort is not sufficient for a mood experience.

This doesn't yet entail the dependence thesis. It is possible that something other than the patterned modification accounts for a mood experience. But considerations analogous to those of Case 1 can be harnessed to show that a mood experience cannot occur in the absence of the patterned modification of other experiences.

Case 2: A mood experience cannot occur in the absence of the patterned modification.

Suppose now that a mood experience can obtain in the absence of a patterned modification. If this is right, then you might be experiencing melancholia (maybe you're having a melancholia raw-feel; or you're representing an unattached affective property *sadness*) while none of your non-mood experiences are affected in any way. However, this entails that it's possible for you to have a melancholy mood experience while nothing seems sad and pointless to you. This is counterintuitive and *prima facie* inconsistent with the moods' global character. But it seems to follow from the view that you don't need a patterned modification of your mental states to undergo a mood experience. Hence, I conclude that a mood experience cannot occur in the absence of a patterned modification of other experiences.

The same conclusion can be reached in a different way, as follows:

Case 3: A mood experience cannot occur in the absence of a patterned modification.

Suppose your experiences are just as they normally are, and yet you feel depressed. So, you could say: ‘nothing seems more bleak and hopeless today than on any other day, but I somehow feel depressed.’ It follows that you could well feel depressed on one day, and elated on another, without *any change* to how things seem to you. So, you could say ‘none of my (non-mood) experiences are in any way different today when compared to yesterday, but I now feel elated, and yesterday I felt depressed.’ This is counterintuitive and *prima facie* inconsistent with the moods’ global character. I doubt that we can make sense of such a radical change in moods without a slightest change in how things seem to the subject. So, this gives us yet another reason to think that a mood experience cannot occur in the absence of a patterned modification of other experiences.

To the extent that the quotations from the previous section articulate an intuitive understanding of what a mood experience is like, the strangeness of Cases 1—3 gives us a good reason to believe that a mood experience depends on a patterned modification of other mood experiences, and that having a basic experience on its own (i.e. without the modifications) is not sufficient.

Moreover, I think that PMV *by itself* explains mood phenomenology better than any alternative view so far presented. So we don’t need basic experiences at all.

5.5 Explaining Mood Phenomenology

In this section, I argue that PMV can account for all the distinctive features of mood phenomenology, while the competing views all face at least one hurdle. This will show that all we need to account for mood phenomenology is a patterned modification. Hence, nothing else is required.

5.5.1 Characteristic Features of Moods

The characteristic features of a mood experience are as follows: mood experiences are (1) global, (2) seemingly non-intentional and (3) phenomenologically similar to emotions.

- I have explained *globality* above in more detail but it is important to recall the caveat that *globality does not have to be total*. You can have a mood experience even when many but *not all* of your mental states are affected.
- *Non-intentionality of moods* is usually expressed by the idea that moods don't typically seem to be about anything in particular. Introspectively, it's often hard to specify what you're depressed or elated about.
- *Similarity between moods and emotions* is usually illustrated by examples: being anxious, and fearing the wolf; being elated, and being happy that the Eagles won; being depressed and being sad that the Eagles lost. All these pairs of experiences seem to share important experiential aspects. They're phenomenally similar in some relevant way.

There are two more characteristics, however, that I'd like to add.

- Sometimes *moods are introspectively obvious*, other times they are *introspectively opaque*. Sometimes, you know you're in a lousy mood. Other times, other people have to point out your mood to you for you to notice it. Let's call this characteristic

the moods' (4) *epistemic duality*. (4) has not gone unnoticed among mood theorists (see e.g. Paul Griffiths (1989)), but it's not mentioned nearly as often as (1)—(3).

- Lastly, it seems that it would be good for any theory of mood experience to *explain why MANIFESTATION is true* or at least gesture at an alternative mechanism of how the moods' global character occurs at the level of consciousness.

Below is a schematic list of mood characteristics, and an indication of where each of the views here considered might have trouble accounting for it (if anywhere). I will first discuss how PMV accounts for all these aspects. Then I will point out some of the shortcomings of its competitors.

Table 5.1 A summary of how various views account for central aspects of mood phenomenology

	View Accounts for:				
	<i>Globality</i>	<i>Apparent Non-intentionality</i>	<i>Similarity to Emotion</i>	<i>MANIFESTATION</i>	<i>Epistemic Duality</i>
Mood Represents Bodily State	No	No	Depends on Theory of Emotion	No	No
Mood Represents World as Whole	No	Yes ¹¹	Yes	Yes	No
Mood Represents Series of Objects	No	No	Yes	Yes	No
Mood Represents Unattached Affective Property	No	Yes	Yes	No	No
Mood Represents One's Overall State	No	Yes	Yes	No	No
Mood is a Raw Feel	No	Yes	Yes	No	No
PMV	Yes	Yes	Yes	Yes	Yes

¹¹ This is debatable: both Mendelovici and Kind appeal to phenomenology to deny plausibility to intentionalist views of this sort. They both hold that, in Mendelovici's words "Some cases of anxiety really don't seem to be plausibly directed at anything at all. One just feels anxious. Similarly, some cases of sudden elation really don't seem to be directed at the world as a whole, an unspecified object, or anything else. One just feels elated. Insisting that these experiences are in fact directed at unusual objects simply gets their phenomenology wrong" (2014). One way to reply to this objection would be simply to deny the phenomenological datum. But a less radical solution is available, at least for 'whole-world' intentionalism. It might be held that to represent something is, perhaps typically or paradigmatically, to discriminate it from the background. But when you represent the world as a whole, it doesn't seem to you that you discriminate anything in particular from the background. Yet, it needn't follow that you don't represent anything at all. Representing without discriminating is atypical of representation in general, but not incoherent (imagine seeing a uniformly colored wall right in front of your eyes).

5.5.2 PMV Explains Everything about Mood Phenomenology

I start with *globality*. Globality, recall, refers to the mood experiences' having wide-reaching effects on other experiences, as well as on the subject's entire psychological life. At the same time, it is not necessary that all of the mood subject's conscious states be somehow affected by the mood. PMV has resources to account for the moods' global character in a non ad-hoc way. This is because a series' being patterned seems to imply at least some globality without implying totality (a pattern can survive some random noise: the series 010100101010101 is still a patterned series).

To see this, let us look at two examples of numerical series: the first one is 0110117485769348766480921. The second is 01101101101101101148576934876. The first, I take it, is not patterned, even though there's a non-random bit there in the beginning. This is because specifying the pattern to the computer (print 011 twice, and then print 748...) would not substantially reduce the amount of information you have to provide the computer with to transmit the series. Things are different with the second series. Here, the opportunity for the compression of the data is bigger. You may instruct the computer to print '011' six times, and then specify verbatim the remaining sequence. You would *substantially* compress the information in this way. The compressing algorithm would be *useful*. But this seems to be because there is a substantial amount of data in the series that can be compressed, and not just a small orderly bit somewhere. In general, it does seem that for us to talk about a patterned series (or whatever), it has to be usefully compressible, and for it to be usefully compressible, the compression must be available for a substantial proportion of the bits comprising the series. Hence, patterns have a global character.

Now, you might think that the use of imprecise terms like 'substantially' should be avoided in what's claimed to be a formal treatment of patterns. But the mathematical

definition of the pattern that I'm using likewise involves some imprecise wording ("about the same"). Similarly, other philosophical explications of the mathematical view of randomness (cf. (Eagle, 2014)) appeal to such informal considerations when determining whether there is a pattern. Ultimately, it looks like talk of "*significantly* longer output strings" and "usefulness" (cf. *ibid.*) is extremely helpful even in more formal settings. It doesn't seem to me that this amount of vagueness in the application of the terms "pattern" (and, if my view is right, "mood") should trouble us. "Mood" is a common sense concept, so it is to be expected that cases will arise where pragmatic considerations will turn out to be the final arbiter as to the term's applicability.

So it appears that PMV can give a theoretically motivated account of a mood experience's global character. While it could ultimately be a matter of practical expediency which algorithms are useful, it appears intuitively right to say that *a substantial part of a series has to be compressible* for us to judge an algorithm to be useful, and hence, a series to be patterned. Similarly, a substantial portion of one's experiences has to be modified for us to judge a person as having a mood experience. The globality of moods is a special case of the globality of patterned series in general.

When it comes to explaining mood experiences' *apparent lack of intentional objects*, PMV can do the job as well. The reason to think so is that patterned series in general do not appear to have intentional objects. Consider a series of numbers, like [1, 3, 6, 10]. Obviously, each numeral refers to a number. But it's not so clear that the whole series is about numbers as well (which numbers? The ones in the series? But the algorithm specifying the pattern allows us to continue the series indefinitely. Is the series about those numbers that would come after 10 were the series to be expanded, too?). Similarly, consider

Dennett's use of barcodes as prime examples of patterns. It's not immediately obvious what, if anything, these barcodes are about.

Now you may think that since its bits are intentional states, the patterned series of affective modifications also has intentional objects: the objects of the particular states. But this doesn't follow. Even though a mood experience consists in a pattern of mental states that do have intentional objects, a patterned series doesn't need to inherit all the properties of its elements. Furthermore, the intentional content of its bits appears *irrelevant*, on the face of it, to whether there is a pattern of *affective* modification. Suppose I see a car, a cat, and a caterpillar, while thinking about China, and experience all of them as scary. It's not their intentional contents that make it the case that these states exhibit a pattern. It's their affective modifications.

To this it might be replied that affective modifications themselves are (or could be) intentional (maybe they're about affective properties things are represented as having). Hence, the patterned series itself is about these affective properties. Again, I think this argument fails. To see this, consider the following list of words: *wicket, goal, striker, match, possession, tactics, midfield, touchdown*.¹² This series of words exhibits a pattern: aside from the first and last one, all the words are connected with soccer. The series exhibits the pattern in virtue of the intentional properties of the words. But the series itself (six soccer words flanked by two non-soccer words) is not about soccer, or anything else in any

¹² Worries might arise as to whether an algorithmic description of this series can be given; that is, whether there can be patterns "at the intentional level," as Dennett (1991) puts it, that conform to the mathematical definition. If there cannot, however, I reiterate my proposal to think of the algorithmic understanding of patterns merely as brining out more clearly the characteristics of patterns and patterned series that our intuitive grasp of these notions already attributes to them (i.e. globality and apparent non-intentionality).

clear or obvious sense. By this I mean that it's hard to decide, on the face of it, whether its object is the game of soccer as a whole, the things its elements refer to, or nothing at all.

The intentionality of most moods is not transparent to introspection. It is usually not obvious, from the inside, whether moods are about anything or not; it is, indeed, a puzzling question on the face of it. The same puzzlement arises when we ask about intentional properties of patterned series in general (it's true even in the apparently easiest cases, like the soccer pattern above). Thus, the lack of an obvious intentional object of a mood looks like a special case of a patterned series lacking an obvious intentional object *even when its elements all have intentional properties and when the series is patterned in virtue of these intentional properties*.

It is worth noting, too, that some moods have fairly clear objects: being irritable, for instance, simply means that things and events you encounter irritate you. Similarly, some patterns clearly are about the items the pattern's elements refer to. Consider, for instance, a shopping list (arranged corresponding to where in the supermarket each item is: first the dairy products, then fruits and vegetables, then soft drinks). It seems clear that, in this case, the list is about the items its elements are about. But just as moods are typically seemingly non-intentional, so the example of a shopping list could be considered atypical of patterned series.

PMV can also account for the *phenomenological similarity* between emotions and moods. Mood phenomenology depends on there being a pattern of affective modifications appropriate for the mood in question. The best way to cash out appropriateness is in terms of phenomenological similarity, roughly as follows: for an affective modification of an experience to be appropriate for the mood in question, it needs to exhibit phenomenological

similarity to the emotional experience corresponding to the mood. Otherwise we couldn't speak of that particular mood experience at all.

To see this more clearly, consider the case of Luke, afraid of the squirrel he sees in front of him. Luke's experience is phenomenally different from just seeing the squirrel without fearing it. In contrast to his neutral experience, Luke's experience of the squirrel as threatening involves some sort of modification by an affective component. Suppose now that the encounter with the squirrel has put Luke in an anxious mood. That is, on the view I'm proposing, his overall experience can now be characterized by many of his non-mood experiences' being modified in the way appropriate for anxiety. That is, in a way that is roughly similar to how his experience of the squirrel that caused the anxiety¹³ was modified (plausibly, unless the anxiety is overwhelming, the modifications composing the mood are less pronounced than the corresponding modification that is part of the emotional experience¹⁴). Hence, there is a phenomenological similarity between the phenomenal character of Luke's emotional experience of the squirrel and the phenomenal character of his anxiety, because the affective components of an emotional experience and the affective modifications that compose the mood experience are phenomenologically similar.

The next characteristic of mood experiences I consider is their epistemic duality. Sometimes, we can easily tell that we're in a mood. Other times, it doesn't occur to us that we are in a mood until some external observer points it out (or we reflect back on our past experiences). PMV has resources to account for this: some patterns are easy to spot, and

¹³ Causal relations obtaining between emotions and moods is an artifact of the example, and by no means essential to my point.

¹⁴ I am not committing myself here to any theory of emotional experience. I'm merely committing myself to the claim that fearing the snake in virtue of seeing it as scary is phenomenally different than just seeing the snake.

others aren't (see, e.g., (Dennett, 1991) again, and (Checkosky & Whitlock, 1973)). Epistemic duality of moods is a special case of the epistemic duality of patterns in general.

Lastly, PMV explains the truth of MANIFESTATION. Mood experiences necessarily involve modifications of non-mood experiences because they are wholly dependent on these modifications.

In total, PMV explains all the characteristic features of the mood experiences in a theoretically satisfying way, either by appealing to features that patterned series have, or by appealing to features that affective modifications have.

5.5.3 Mood Phenomenology and the Competitors to PMVs

In this section I argue that each of the competitors to PMV faces hurdles that my view does not in accounting for the characteristic features of moods. I will only here justify the negative entries in the table above. I will assume without argument that the positive entries about the competing views are correct.

It appears that each competitor view fares worse than PMV because of their common failure to account for the moods' epistemic duality. There is nothing about the raw feels in general, or about the representational states in general, that renders some of them easily noticeable and others hard to notice. We must therefore treat it as a brute fact that some mood experiences are, and others aren't easily accessible by introspection. This is in marked contrast to PMV. I have mentioned, however, that epistemic ambiguity is not a universally agreed-upon characteristic of moods. I will bracket it then in what follows.

I will start with the views that mood experience represents one's overall state or some state of one's bodily equilibrium. The problem for these views is that if it's the subject's overall state that gets represented in a mood experience, then you have to have a specific

theory of emotion to explain the *phenomenological similarity* between the two without any *ad hoc* devices. In particular, your theory of emotional experience must postulate an essential connection between such an experience and the bodily experiences that Tye's brief description lists (alternatively, it must postulate an essential connection between your experience of emotion and your assessment of your overall state). This already indicates that the view is in a worse position than the remaining ones to account for mood phenomenology, because others are not committed to any particular account of the phenomenology of emotions.

Additionally, on the one hand, while postulating an essential connection between bodily experience and emotional experience is an attractive position, the bodily-representation view of moods makes it hard to see why mood experiences would be global; it's not clear what it is about my representing my body that should influence my other perceptions, thoughts and feelings. Conscious states of awareness of our bodies don't generally do that. Right now, I'm aware of myself as sitting on a chair, with my legs uncrossed. This doesn't have some global influence on my perceptions, or thoughts. On the other hand, while it could be that a representation of one's overall state is in some way pervasive, its similarity to emotional experiences is harder to argue for (why would emotional experiences involve an assessment of one's overall state, exactly?).

In Error! Reference source not found. I have indicated that whole-world and series-of-objects versions of representationalism cannot account for the mood experiences' globality. Initially, this seems odd. After all, when you represent *everything* in one representational state, you get the global character of that state more or less 'for free.' A state that represents everything is, perhaps by definition, global. The worry here, however,

is that such a global character is also *total*. The mood experience, on these views, would encompass *everything* that the subject represents. But, as I have argued before, globality should not be taken to entail that *all* of your conscious states contribute to the mood phenomenology. They seem to, however, on these intentionalist accounts. The way mood experiences manifest themselves in consciousness must, on these views, involve *all* of the subject's non-mood experiences. Having this implausible consequence is a disadvantage of whole-world and series-of-objects Intentionalism.

On the other hand, if proponents of these views were to claim that only a portion of phenomenal states needs be affected for the subject to be in a mood, there would be no theoretical justification for this relaxation of the conditions for the occurrence of a mood experience. The relaxation would have to remain unexplained. This is an explanatory debt on the part of these theories, one that PMV doesn't incur.

When it comes to Mendelovici's account of moods as representing unattached affective properties the problem also arises for its ability to explain the *globality* of moods. Just because a property you're experiencing is not experienced as a property of some particular thing, it doesn't follow it would pervade your entire experience. This is because, in general, experiencing an unattached property does not seem to confer globality on the experience. Consider other putative examples of experiencing an unattached property: "color experience in a ganzfeld, motion experience in peripheral vision, or experiences had in situations where binding failures are prone to happen [as is the case in patients with Balint's syndrome]" (Lee, 2014), or experiences of unbound motion that Type-2 blindsight patients report (cf. (Weiskrantz, 2002). I'm not going to pass judgment on whether these really are unbound features (and neither does Lee). However, it is not clear that experiencing any of

these potentially unattached properties automatically exerts global effects on other experiences the subject is also having (even the ganzfeld experience, despite its name, need not have any effects on my thoughts and feelings, or on what I'm hearing or tasting, for that matter). So, the mere fact that a property is represented as unattached does not make the representation a good candidate for globality.

The raw feel view faces roughly the same difficulty that plagues Mendelovici's account; it is at a loss when it comes to accounting for a mood experience's globality. While moods' globality is sometimes *explained in terms of* their non-intentionality (see, e.g., Siemer (2009) and Frijda (1994): moods, so the thought goes, seem global *because* they are not confined to being about some single intentional object), this explanation is inadequate.

This is because globality does not seem like a general characteristic of putatively non-intentional states. Think about an experience of a phosphene. Arguably, such an experience doesn't have an object. But no one would want to say that it thereby acquires some global influence on your consciousness and behavior, on "the world and everything in it." And if, as some philosophers think, pains, itches and tickles are not intentional, they still do not seem *thereby* to exert some global influence on your experiences. So a conscious state's lacking an intentional object does not explain that state's having some sort of global influence on experience.

It appears then that none of the competing views is able to account for all the features of mood experiences in a principled way. This makes them inferior to PMV.

5.5.4 Amy Kind's Challenge

Amy Kind (2014) has recently offered a challenge to Intentionalism about moods that might also put pressure on PMV. According to Kind, these four claims are perfectly coherent:

[1] My depression hasn't worsened since yesterday, but today everything seems blacker.

[2] My fearfulness hasn't lessened any, but the world seems less scary today.

[3] The world doesn't seem any more worrying today, yet somehow my anxiety has intensified.

[4] The world doesn't seem any blacker today, yet I find myself sinking deeper into depression (2014)

The coherence of the above sentences is meant to show that we could vary intentional content of mood experiences without a change in the phenomenal character of the experience and vice versa. Even though Intentionalism is Kind's main target, PMV should also be able to meet this challenge. This is because PMV also postulates a tight connection between a mood experience and how things seem to the mood subject.

One way to meet Kind's challenge is to explain the coherence of [1]-[4] by claiming that the first conjunct in [1] and [2] and the second conjunct in [3] and [4] are *not* based exclusively on the subject's phenomenology; rather, they're judgments about the subject's overall state based on more than what the world seems like to him. For example, suppose a super-neuroscientist looks at my brain and assures me that the neural correlates of my mood have not changed any of their physical parameters since yesterday. Being a thoroughgoing materialist, I acquire a justified belief that my depression hasn't worsened since yesterday, even though my phenomenology changed. So I may say [1] perfectly coherently without thereby impugning Intentionalism. In general, the judgments about my

own mood state can be based on phenomenally unconscious background beliefs about me (in addition to what my experience feels like), or indeed about the evidential status of my phenomenology as it pertains to my overall state.

Another way to respond to Kind's challenge is to say that the coherence of [1]-[4] comes from the fact that the second conjunct in [1] and [2], and the first conjunct in [3] and [4] is only a partial description of a subject's mood phenomenology (perhaps "everything" could be read as "everything I perceive"). Perhaps everything seems blacker, but I'm no longer entertaining so many sad memories as I did yesterday. Perhaps the world doesn't seem any blacker today, but thoughts about death and suffering just occur more frequently to me. This, I take it, would account for the coherence of [1]-[4] without endangering Intentionalism about moods.

Kind could retort that in proposing these two ways of accounting for the coherence of [1]-[4] I have not taken the reports "at face value." But I think that, when taken at face value (i.e., I assume, as being exclusively concerned with the subject's phenomenology: "My experience of depression has not changed, but the world seems blacker today"), the four reports *do* seem extremely puzzling. Having heard one of those, I think it would be reasonable to ask the speaker for clarification. Some plausible clarifications are offered above.

At any rate, to the extent that Kind's challenge is a problem for Intentionalism, it is also a problem for PMV. Hence, the replies I have offered can serve to dispel the worries Kind' challenge raises for my view as well.

5.6 Mood Experiences and Representationalism

Part of the philosophical interest in mood experiences is due to their being an apparent counterexample to Intentionalism. If I'm right, however, then mood experiences do not provide much insight into that issue. That's because mood experiences, in my view, cannot be cited as straight-up counterexamples to Intentionalism nor as straight-up refutations of anti-intentionalist concerns. This is because we need to decide both on (a) how modification works (is the modification itself intentional, like a seeming; or is it non-intentional, like an adverbial modification?) *and* (b) which properties of the elements get inherited by the patterned series, in order to be able to decide whether mood-experiences really are a problem for Intentionalism. Appealing to mood phenomenology in that debate without considering these issues first is to put the cart before the horse. That's because mood phenomenology is compatible with all these options (indeed, as I have mentioned in a footnote above, the apparent non-intentionality of moods is also compatible with some explicitly intentionalist accounts).

5.7 Mood Experiences and Spatial Unity of Consciousness

In this section, I argue that this account of mood experiences is compatible with the account of phenomenal unity of consciousness in terms of spatial unity. Elizabeth Schechter (2013) distinguishes two kinds of the unity of consciousness: coherence unity and phenomenal unity. Phenomenal unity (also called "co-consciousness") is characterized as follows: "for two experiences to be co-conscious with each other is simply for there to be something it is like for their subject to undergo both of those experiences simultaneously—something that is different from what it is like to undergo the first experience one day and then the second the next, for instance, or something that is different

from what it is like for Joan to undergo the first experience while, simultaneously, John undergoes the second.”

Coherence unity, on the other hand, obtains when there is integration between contents of the subject’s states. According to Schechter, “[a] subject with a coherence unified consciousness possesses, first of all, conscious experiences that make sense together, by presenting a unified world, a world that makes sense.” It’s worth noting that Schechter’s understanding of coherence is more demanding than the one I presented in the Introduction. In fact, it’s more akin to Shoemaker’s representational unity.

There is another kind of unity that, I think, deserves to be called “phenomenal.” Think, for a moment, about experiences in a weak-intentionalist fashion. On such a view, experiences are composed of the representational content and its mode of presentation. However, the way in which the content and the mode of presentation are connected to yield a single experience is not the same as the way in which John’s gustatory experience and his auditory experience are parts of some bigger experience that John is enjoying. Similarly, think about Murat Aydede’s (forthcoming) idea that some experiences are modified by an affective quality. This view would also have it that the modified experience consists of the content and the way in which that content is presented. This kind of unity, between the “what” and the “how” of an experience is definitely different than the kind of unity that obtains between John’s two experiences. It is not my purpose to say in virtue of what this happens. I’m pointing out to there being such an “intra-experiential” unity, in addition to the kinds of unity normally discussed under the heading of the unity of consciousness.

Now, it seems plausible that an account of phenomenal unity would *not* be, primarily, an account of intra-experiential unity. The way in which the mode of presentation and the content of an experience are unified is not the way in which two distinct experiences are unified. The basic point is that experiences are in some way free-standing entities, ones that can, but don't have to, stand in phenomenal unity relations to others. Modifications are ontologically dependent entities. Even those who think that unified consciousness consists in just one experience, and those who think that an experience's phenomenal character depends on what other experiences it is unified with, should, I think, be able to make sense of this distinction. They should have in their theory a place for the distinction between what I here call "intra-experiential" and "inter-experiential" unity, even though, of course, these labels would not be, strictly speaking, appropriate to use for them.

Because of what I have mentioned above, the following line of criticism of my view of phenomenal unity cannot get off the ground: "your account says that spatial unity is necessary for phenomenal unity; but a modification of an experience is unified with what it modifies without there being any spatial unity relation obtaining between them. Hence, spatial unity is not necessary for phenomenal unity." It cannot get off the ground because it equivocates between two senses of phenomenal unity: the inter-experiential unity and the intra-experiential unity. We should not expect a single account to capture both kinds of unity. And neither do those authors who advance theories of unity. The account that is given is the account of inter-experiential phenomenal unity, not an account of intra-experiential unity.

Because of this, the idea that moods wholly depend on patterns of modifications of experiences is compatible with thinking that spatial unity is necessary for (inter-

experiential) phenomenal unity. Such modifications are not free-standing experiences for which the question of unity arises. This is because the unity relation is a relation between experiences, not between non-detachable aspects of a single, free-standing experience. Hence, for the mood experience to obtain, other experiences must be simply be modified in certain ways. The kind of connection between an experience and its modification is not the connection that unity of consciousness theorists seem interested in. On the other hand, the *experiences* that collectively give rise to the mood experience are spatially unified. Hence, mood experiences do not pose a special problem for my account.

Even supposing that the modifications can best be accounted for in terms of seemings, it looks like seemings too are not self-standing parts of an experience. You can't have a conscious *perceptual* seeming without a perception. So seemings too exhibit a sort of intra-experiential unity with their corresponding perceptions.

This shows that my account of the phenomenal character of moods in terms of modifications to non-mood experiences is still compatible with the spatial unity thesis.

5.8 Summary

In this chapter I have argued that a patterned modification of phenomenal characters of a subject's mental states explains everything we need to explain about mood phenomenology. Then I argued that while apparent non-intentionality of moods is compatible with a variety of accounts (including intentionalist ones), the competing views I have considered cannot explain the special features of mood experiences without employing some ad hoc devices. I noted the relevance of this paper for the appropriateness of the continued use of mood experiences as counterexamples to Intentionalism. Lastly, I explained the compatibility of this account of moods with the spatiality thesis.

Chapter 6: Conclusion

In this dissertation I have defended the view that unity of consciousness requires unity of conscious spatial representation. There are two principal reasons to doubt the necessity claim: first, some phenomenally conscious states might not be spatial at all. Examples include moods and conscious thoughts. Second, it appears that there could be spatial experiences that are both phenomenally unified and spatially disunified. My approach has been to develop my theory with a view to laying both these types of objection to rest.

I have done so by first proposing the thesis that our perceptual representation of space is always a representation of a single, unified space. I then used the thesis to defend the necessity of spatial unity against objections from Bayne, Dainton, Tye, Roelofs, and Watzl; the objections are to the effect that two phenomenally unified spatial experiences can fail to be spatially unified. This work was done in Chapter 2.

In Chapter 3 I have first extended spatial unity to non-perceptual spatial experiences: imagining, recollecting, and enjoying afterimages and phosphenes. In the second part of this chapter, I have also considered and rejected the position that some perceptual experiences of objects are not experiences as of objects in space. Together, Chapters 2 and 3 show that spatial unity obtains across all experiences standardly taken to have spatial phenomenology.

In Chapter 4 I have argued that phenomenally conscious thoughts should be assumed to have spatial phenomenology in the absence of strong arguments to the contrary. I have outlined reasons why accepting spatial cognitive phenomenology should be a default position; then I have looked at some objections to granting spatial phenomenology to

thoughts, and found them all wanting. Lastly, I have indicated how the space represented when having thoughts is represented as related to the space we perceive.

In Chapter 5 I have taken on another case of a phenomenally conscious state (mood) that might be held to lack spatial phenomenology. I have argued that mood phenomenology is nothing over and above the phenomenology of other experiences suitably modified in a patterned way. From this it follows that mood phenomenology, in virtue of being wholly derived, does not present an additional challenge to my overall thesis. Mood phenomenology depends on phenomenally conscious perceptual, cognitive and somatic states; since these states have spatial phenomenology, and since they exhibit spatial unity, mood phenomenology cannot obtain in the absence of spatial unity.

I take the above chapters to establish, then, that (1) all experiences are spatial experiences or wholly depend for their phenomenology on spatial experiences, and (2) that all spatial experiences represent the same space. From this, it follows that there cannot be any two phenomenally unified experiences that are not spatially unified. Spatial unity is necessary for phenomenal unity.

I have thus vindicated a naively appealing view of phenomenal unity. I have done so in terms of representations whose neuro-functional profiles are relatively well-understood. I have preserved the idea that unified consciousness has something to do with a picture of a unified world, without overintellectualizing the resources required to build up that picture. The view I'm proposing also explains why unified consciousness frequently has perspectival structure, while allowing for the possibility that such perspectival structure is not necessary for unified consciousness (because it's not necessary for conscious spatial representation). Lastly, I have also outlined a novel theory of mood experiences, one that

allows us to see more clearly what role moods play (and should play) in the debate over Intentionalism.

My dissertation provides a first step in an account of the structure of consciousness. I envisage my next step to consist in investigating further the spatial structure that consciousness necessarily displays, both in its ordinary form as well as in its disruptions. In particular, I want to explore further the nature of perspectival representation; there is some temptation to think of unified consciousness as what constitutes our conscious *perspective* on the world. There is also some temptation to think of spatial representation as essentially perspectival. And work is being done in both philosophy and neuroscience to explain consciousness in terms of perspectival (egocentric) representations. Unfortunately, these attempts seem to run into problems if my view is right: conscious spatial unity does not seem to demand perspectivalness. But since conscious spatial unity is not essentially perspectival, and since it underlies the structure of unified consciousness, it appears that we should not attempt to explain consciousness in terms of perspectival representation; neither should we seek to locate the perspectival character of unified consciousness in spatial unity alone. Two different lines of thought suggest themselves here: either the understanding of unified consciousness as necessarily having a perspectival structure, a structure with a center and a periphery, is a matter of some additional aspect of our phenomenology (attention perhaps); or, alternatively, unified consciousness doesn't *have to* have a perspectival structure; that it typically does exhibit perspectival character, on the other hand, may well be due to its having a spatial structure that is also frequently a matter of perspectival representation. This is certainly an interesting project to inquire into.

Another area into which the considerations of this dissertation lead concerns the prospects and viability of providing a full structural account of the unity of consciousness; that is, providing necessary and sufficient conditions for consciousness to be unified. We may wonder, in particular: is anything else necessary for a unified consciousness? What is sufficient for consciousness to be unified? My primary aim in this area of investigation will be to show that an account of sufficient conditions for unified consciousness is either impossible or trivial. Hence, those who seek to provide a structural account of consciousness should focus on necessary conditions for unified consciousness.

While my view is inconsistent with other views on the unity of consciousness insofar as those other views deny the necessity of spatial unity, it is not completely at odds with them. And so, my view simply grants consciousness more structure than Bayne's mereological view, Tye's one-experience view, and Dainton's view of co-consciousness as a primitive relation. On my view, consciousness has a more uniform structure than on Masrour's Leibnizian view. Intriguing prospects arise for connecting the view I defend with Watzl's position according to which the structure of consciousness is given by an attentional field, as well as similar ideas present in Roelofs's and Chudnoff's work. I think it is also worth exploring what my view has to offer to Shoemaker's conception of the unity of consciousness.

Lastly, an important consequence of the view I'm proposing is the severance of the link between spatial awareness and awareness of objectivity. To the extent that spatial awareness displays a certain degree of uniformity across different kinds of experiences (e.g. experiences of objects, images, phosphenes, pains, thoughts, and the like), that is, to

the extent that these experiences display *spatial* unity, then representing something as in space cannot by itself constitute representing it as a denizen of mind-independent reality.

The results I report here, then, seem to have some application both to theorizing about the unity of consciousness, and for broader issues in philosophy of mind that do not directly connect with such theorizing. I think this is quite interesting and the opportunities for expanding this project's scope are plentiful.

Bibliography

- Armstrong, K. C., & Ramachandran, V. S. (2003). Projecting sensations to external objects: Evidence from skin conductance response. *Proceedings of the Royal Society of London B: Biological Sciences*, 270(1523), 1499-1506.
- Aydede, M. (forthcoming). A contemporary account of sensory pleasure. In L. Shapiro (Ed.), *Pleasure: a History*. Oxford: Oxford University Press.
- Batty, C. (2010). A representational account of olfactory experience. *Canadian Journal of Philosophy*, 40(4), 511-538.
- Baylis, G. C., & Baylis, L. L. (2001). Visually misguided reaching in Balint's syndrome. *Neuropsychologia*, 39(8), 865-875.
- Bayne, T. J. (2004). Self-consciousness and the unity of consciousness. *The Monist*, 87(2), 219-236.
- Bayne, T. J. (2010). *The unity of consciousness*. Oxford, England; New York: Oxford University Press.
- Bayne, T. J., & Chalmers, D. J. (2003). What is the unity of consciousness? In A. Cleeremans (Ed.), *The Unity of Consciousness*. Oxford: Oxford University Press.
- Bermudez, J. L. (2003). *Thinking without words*. Oxford: Oxford University Press.
- Bortolotti, L., & Broome, M. (2009). A role for ownership and authorship in the analysis of thought insertion. *Phenomenology and the Cognitive Sciences*, 8(2), 205-224.
- Brogaard, B. B. (2012). Vision for action and the contents of perception. *The Journal of Philosophy*, 109(10), 569-587.
- Byrne, A. (2010). Recollection, perception, imagination. *Philosophical Studies*, 148(1), 15-26.
- Campbell, J. (1993). The role of physical objects in spatial thinking. In N. Eilan, R. McCarthy, & B. Brewer (Eds.), *Spatial representation: Problems in philosophy and psychology* (pp. 65-95). Oxford; Cambridge Mass.: Blackwell.
- Carruthers, P., & Veillet, B. (2011). The case against cognitive phenomenology. In T. Bayne & M. Montague (Eds.), *Cognitive phenomenology* (pp. 35-56). Oxford: Oxford University Press.

- Chaitin, G. J. (1975). Randomness and mathematical proof. *Scientific American*, 232(5), 47-52.
- Chalmers, D. (2006). Perception and the fall from Eden. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 49-125). Oxford: Oxford University Press.
- Checkosky, S. F., & Whitlock, D. (1973). Effects of pattern goodness on recognition time in a memory search task. *Journal of Experimental Psychology*, 100(2), 341-348.
- Chen, L. (1982). Topological structure in visual perception. *Science*, 218(4573), 699-700.
- Chudnoff, E. (2013). Gurwitsch's phenomenal holism. *Phenomenology and the Cognitive Sciences*, 12(3), 559-578.
- Chudnoff, E. (2015). *Cognitive phenomenology*. London; New York: Routledge.
- Connolly, K. (2014). Perceptual learning and the contents of perception. *Erkenntnis*, 79(6), 1407-1418.
- Dainton, B. (2000). *Stream of consciousness: Unity and continuity in conscious experience*. London ; New York: Routledge.
- Dainton, B. (2003). Unity in the void: Reply to Revonsuo. *Psyche*, 9(8).
- Dennett, D. (1981). Where am I? . In D. Dennett & D. Hofstadter (Eds.), *The mind's I: Fantasies and reflections on mind and soul* (pp. 217-231). New York: Basic Books.
- Dennett, D. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27-51.
- Dixon, M. J., Smilek, D., & Merikle, P. M. (2004). Not all synaesthetes are created equal: Projector versus associator synaesthetes. *Cognitive, Affective, & Behavioral Neuroscience*, 4(3), 335-343.
- Dokic, J., & Pacherie, E. (2006). On the very idea of a frame of reference. In M. Hickmann & S. Robert (Eds.), *Space in Languages: Linguistic systems and cognitive categories* (pp. 259-280). Amsterdam: John Benjamins.
- Dummett, M. (1994). *Origins of analytical philosophy*. Cambridge, Mass.: Harvard University Press.
- Eagle, A. (2014). Chance versus randomness. *The Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/archives/spr2014/entries/chance-randomness>
- Eilan, N., McCarthy, R. A., & Brewer, B. (1993). *Spatial representation: Problems in philosophy and psychology*. Oxford England; Cambridge, Mass.: Blackwell.
- Evans, G. (1982). *The varieties of reference*. New York: Oxford University Press.

- Evans, G. (1985). Molyneux's question. In G. Evans (Ed.), *Collected Papers*: Oxford University Press.
- Firestone, C. (2013). How “paternalistic” is spatial perception? Why wearing a heavy backpack doesn’t—and couldn’t—make hills look steeper. *Perspectives on Psychological Science*, 8(4), 455-473.
- Fish, W. (2005). Emotions, moods, and intentionality. In G. Forrai & G. Kampis (Eds.), *Intentionality: Past and future* (pp. 25-36). Amsterdam; New York: Rodopi.
- Flegg, G. (1974). *From geometry to topology*. London: English Universities Press.
- Friedman-Hill, S., Robertson, L. C., & Treisman, A. (1995). Parietal contributions to visual feature-binding: Evidence from a patient with bilateral lesions. *Science*, 269, 853-855.
- Frijda, N. (1994). Varieties of affect: Emotions and episodes, moods, and sentiments. In P. Ekman & R. Davidson (Eds.), *The nature of emotion* (pp. 59-67). Oxford; New York: Oxford University Press.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in neurosciences*, 15(1), 20-25.
- Green, E. J. (forthcoming). A layered view of shape perception. *British Journal for the Philosophy of Science*.
- Gregory, D. (2010). Imagery, the imagination and experience. *Philosophical Quarterly*, 60(241), 735-753.
- Gregory, D. (2013). *Showing, sensing, and seeming: Distinctively sensory representations and their contents*. Oxford: Oxford University Press.
- Griffiths, P. E. (1989). Folk, functional and neurochemical aspects of mood. *Philosophical Psychology*, 2(1), 17-30.
- Haugeland, J. (1978). The nature and plausibility of cognitivism. *Behavioral and Brain Sciences*, 1(2), 215-226.
- Hurlburt, R. T. (1997). Randomly sampling thinking in the natural environment. *Journal of Consulting and Clinical Psychology*, 65(6), 941.
- Hurlburt, R. T., & Akhter, S. A. (2008). Unsymbolized thinking. *Consciousness and Cognition*, 17(4), 1364-1374.
- Hurley, S. (1998). *Consciousness in action*. Cambridge, Mass.: Harvard University Press.

- Husserl, E. (1913/1952). *Ideas pertaining to a pure phenomenology and to a phenomenological philosophy* (R. Rojcewicz & A. Schuwer, Trans.). Dordrecht: Kluwer Academic Publishers.
- Kelly, S. D. (2010). The normative nature of perceptual experience. In B. Nanay (Ed.), *Perceiving the World* (pp. 146-160). Oxford: Oxford University Press.
- Kenny, A. (1963). *Action, emotion and will*. London,: Routledge & K. Paul.
- Kind, A. (2014). The case against representationalism about moods. In U. Kriegel (Ed.), *Current controversies in philosophy of mind* (pp. 113-134). New York; London: Routledge.
- Klatzky, R. (1998). Allocentric and egocentric spatial representations: Definitions, distinctions, and interconnections. In C. Freksa, C. Habel, & K. Wender (Eds.), *Spatial cognition* (Vol. 1404, pp. 1-17): Springer Berlin Heidelberg.
- Koehler, K. (1979). First rank symptoms of schizophrenia: Questions concerning clinical boundaries. *British Journal of Psychiatry*, 134, 236-248.
- Kriegel, U. (forthcoming). *The varieties of consciousness*. Oxford: Oxford University Press.
- Lee, G. (2014). Experiences and their parts. In D. Bennett & C. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 287-322). Cambridge, Mas.: The MIT Press.
- Levinson, S. C. (1996). Frames of reference and Molyneux's question: Crosslinguistic evidence. In P. Bloom, L. Nadel, M. A. Peterson, & M. F. Garrett (Eds.), *Language and space* (pp. 109-169). Cambridge, Mass.: MIT Press.
- Lloyd, D. M. (2007). Spatial limits on referred touch to an alien limb may reflect boundaries of visuo-tactile peripersonal space surrounding the hand. *Brain and cognition*, 64(1), 104-109.
- Lormand, E. (1985). Toward a theory of moods. *Philosophical Studies*, 47(3), 385-407.
- Martin, J.-R., & Pacherie, E. (2013). Out of nowhere: Thought insertion, ownership and context-integration. *Consciousness and Cognition*, 22(1), 111-122.
- Martin, M. G. F. (2002). The transparency of experience. *Mind & Language*, 17, 376-425.
- Masrour, F. (2014). Unity of consciousness: Advertisement for a Leibnizian view. In D. Bennett & C. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 323-346). Cambridge, Mass.: MIT Press.

- Matthen, M. (2005). *Seeing, doing, and knowing: A philosophical theory of sense perception*. Oxford; New York: Oxford University Press.
- Matthen, M. (2014). Image content. In B. Brogaard (Ed.), *Does perception have content?* (pp. 265-290). Oxford; New York: Oxford University Press.
- Mellor, C. S. (1970). First rank symptoms of schizophrenia. *British Journal of Psychiatry*, *117*, 15-23.
- Mendelovici, A. (2014). Pure intentionalism about moods and emotions. In U. Kriegel (Ed.), *Current controversies in philosophy of mind*. London; New York: Routledge.
- Merleau-Ponty, M. (1998). *Phenomenology of perception* (C. Smith, Trans.). London: Routledge.
- Morris, W., & Schnurr, P. (1989). *Mood: The frame of mind*. New York: Springer.
- O'Grady, J. (1990). The prevalence and diagnostic significance of Schneiderian first-rank symptoms in a random sample of acute psychiatric in-patients. *British Journal of Psychiatry*, *156*, 496-500.
- O'Callaghan, C. (2010). Perceiving the locations of sounds. *Review of Philosophy and Psychology*, *1*(1), 123-140.
- Parnas, J., Moller, P., Kircher, T., Thalbitzer, J., Jansson, L., Handest, P., & Zahavi, D. (2005). EASE: Examination of anomalous self-experience. *Psychopathology*, *38*(5), 236-258.
- Pautz, A. (2013). Does phenomenology ground mental content? In U. Kriegel (Ed.), *Phenomenal intentionality* (pp. 194-234). Oxford: Oxford University Press.
- Pawar, A. V., & Spence, S. A. (2003). Defining thought broadcast: Semi-structured literature review. *The British Journal of Psychiatry*, *183*(4), 287-291.
- Peacocke, C. (1985). Imagination, experience, and possibility: A Berkeleian view defended. In J. Foster & H. Robinson (Eds.), *Essays on Berkeley* (pp. 19-35). Oxford: Clarendon Press.
- Peacocke, C. (1992). *A study of concepts*. Cambridge, Mass: MIT Press.
- Peacocke, C. (2014). *The mirror of the world: Subjects, consciousness, and self-consciousness*. Oxford: Oxford University Press.
- Piaget, J., & Inhelder, B. (1963). *The child's conception of space* (F. Lagdon & J. Lunzer, Trans.): Routledge & Kegan Paul.

- Pick Jr, H. L., & Lockman, J. J. (1981). From frames of reference to spatial representations. In L. S. Liben, A. H. Patterson, & N. Newcombe (Eds.), *Spatial representation and behavior across the life span: Theory and application* (pp. 39-61). New York: Academic Press.
- Price, C. (2006). Affect without object: Moods and objectless emotions. *European Journal of Analytic Philosophy*, 2(1), 49-68.
- Prinz, J. (2011). The sensory basis of cognitive phenomenology. In T. Bayne & M. Montague (Eds.), *Cognitive phenomenology*. Oxford: Oxford University Press.
- Rice, H. J., & Rubin, D. C. (2011). Remembering from any angle: The flexibility of visual perspective during retrieval. *Consciousness and Cognition*, 20(3), 568-577.
- Robertson, L. C., Treisman, A., Friedman-Hill, S., & Grabowecky, M. (1997). The interaction of spatial and object pathways: Evidence from Balint's syndrome. *Journal of Cognitive Neuroscience*, 9(3), 295-317.
- Roelofs, L. (2014). What are the dimensions of the conscious field? *Journal of Consciousness Studies*, 21(7-8), 88-104.
- Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, 110(1), 145-172.
- Rusting, C. L. (1998). Personality, mood, and cognitive processing of emotional information: Three conceptual frameworks. *Psychological Bulletin*, 124(2), 165-196.
- Sally, S. L., & Gurnsey, R. (2003). Orientation discrimination in foveal and extra-foveal vision: Effects of stimulus bandwidth and contrast. *Vision Research*, 43(12), 1375-1385.
- Schechter, E. (2013). Two unities of consciousness. *European Journal of Philosophy*, 21(2), 197-218.
- Schwenkler, J. (2012). Does visual spatial awareness require the visual awareness of space. *Mind & Language*, 27(3), 308-329.
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge; New York: Cambridge University Press.
- Shoemaker, S. (1996). Unity of consciousness and consciousness of unity *The first-person perspective and other essays*: Cambridge University Press.
- Siegel, S. (2006). Subject and object in the contents of visual experience. *Philosophical Review*, 115, 355-388.
- Siegel, S. (2011). *The contents of visual experience*: Oxford University Press.

- Siemer, M. (2009). Mood experience: Implications of a dispositional theory of moods. *Emotion Review*, 1(3), 256-263.
- Siewert, C. (2011). Phenomenal thought. In T. J. Bayne & M. Montague (Eds.), *Cognitive phenomenology* (pp. 236-267). Oxford: Oxford University Press.
- Smart, J. J. (1959). Sensations and brain processes. *The Philosophical Review*, 141-156.
- Smith, J. (2014). Egocentric space. *International Journal of Philosophical Studies*, 22(3), 409-433.
- Solomon, R. C. (1993). *The passions: Emotions and the meaning of life*. Indianapolis: Hackett Pub. Co.
- Strawson, P. F. (1959). *Individuals, an essay in descriptive metaphysics*. London,: Methuen.
- Taylor, M. A., & Heiser, J. F. (1971). Phenomenology: An alternative approach to diagnosis of mental disease. *Comprehensive Psychiatry*, 12(5), 480-486.
- Thorpe, S. J., Gegenfurtner, K. R., Fabre-Thorpe, M., & Buelthoff, H. H. (2001). Detection of animals in natural images using far peripheral vision. *European Journal of Neuroscience*, 14(5), 869-876.
- Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. Cambridge, Mass.: MIT Press.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, Mass.: MIT Press.
- Tye, M. (2003). *Consciousness and persons: Unity and identity*. Cambridge, Mass.: MIT Press.
- Tye, M. (2008). The experience of emotion: An intentionalist theory. *Revue Internationale de Philosophie*, 25-50.
- Tye, M., & Wright, B. (2011). Is there a phenomenology of thought. In T. Bayne & M. Montague (Eds.), *Cognitive phenomenology*. Oxford: Oxford University Press.
- Van Cleve, J. (2006). Touch, sound, and things without the mind. *Metaphilosophy*, 37(2), 162-182.
- Van Gulick, R. (2014). E pluribus unum: Rethinking the unity of consciousness. In D. Bennett & C. Hill (Eds.), *Sensory Integration and the Unity of Consciousness* (pp. 375-395).
- Watzl, S. (2014). Attentional organization and the unity of consciousness. *Journal of Consciousness Studies*, 21(7-8), 56-87.

Weiskrantz, L. (2002). Prime-sight and blindsight. *Consciousness and Cognition*, 11(4), 568-581.

Williams, B. (1973). Imagination and the self *Problems of the self*. Cambridge: Cambridge University Press.

